

Putting Logic-Based Distributed Systems on Stable Grounds

Tom J. Ameloot*, Jan Van den Bussche
Hasselt University & transnational University of Limburg

William R. Marczak
University of California, Berkeley

Peter Alvaro
University of California, Santa Cruz

Joseph M. Hellerstein
University of California, Berkeley

submitted 5 September 2012; revised 12 November 2013, 24 April 2015; accepted 16 July 2015

Abstract

In the Declarative Networking paradigm, Datalog-like languages are used to express distributed computations. Whereas recently formal operational semantics for these languages have been developed, a corresponding declarative semantics has been lacking so far. The challenge is to capture precisely the amount of nondeterminism that is inherent to distributed computations due to concurrency, networking delays, and asynchronous communication. This paper shows how a declarative, model-based semantics can be obtained by simply using the well-known stable model semantics for Datalog with negation. We show that the model-based semantics matches previously proposed formal operational semantics.

To appear in Theory and Practice of Logic Programming (TPLP).

KEYWORDS: Dedalus, Datalog, stable model semantics, distributed system, asynchronous communication

1 Introduction

Cloud environments have emerged as a modern way to store and manipulate data (Zhang et al. 2010; Cavage 2013). For our purposes, a cloud is a distributed system that should produce output as the result of some computation. We use the common term “node” as a synonym for an individual computer or server in a network.

In recent years, logic programming has been proposed as an attractive foundation for distributed and cloud programming, building on work in declarative networking

* T.J. Ameloot is a Postdoctoral Fellow of the Research Foundation – Flanders (FWO).

(Loo et al. 2009). The essential idea in declarative networking, is that the programmer uses a high-level declarative language (like Datalog) to specify only what has to happen, and not exactly how. For example, the programmer could specify only that certain messages are generated in reply to other messages; the exact technical details to send (and possibly resend) messages over transmission protocols are filled in by some runtime engine. This frees the programmer from thinking in low-level terms that distract from the actual meaning of the specific program at hand. In particular, complex distributed algorithms and protocols can be expressed in relatively few lines of code (Jim 2001; Alvaro et al. 2009; Hellerstein 2010b). Besides the interest in declarative networking, we are also seeing a more general resurgence of Datalog (with negation) (de Moor et al. 2011; Huang et al. 2011). Moreover, issues related to data-oriented distributed computing are receiving attention at database theory conferences (Hellerstein 2010a; Ameloot et al. 2011; Abiteboul et al. 2011; Ameloot and Van den Bussche 2012; Zinn et al. 2012).

One of the latest languages proposed in declarative networking is Dedalus (Alvaro et al. 2009; Alvaro et al. 2011; Hellerstein 2010b), a Datalog-inspired language that has influenced other recent language designs for distributed and cloud computing such as Webdamlog (Abiteboul et al. 2011) and Bloom (Alvaro et al. 2011).

Model-based semantics In this paper, we describe the meaning of distributed Datalog programs using a model-based semantics. This approach contrasts with most previous work in declarative networking, where the meaning of programs was typically described with an operational semantics (Deutsch et al. 2006; Navarro and Rybalchenko 2009; Grumbach and Wang 2010; Ameloot et al. 2011), with a few exceptions (Lobo et al. 2012; Ma et al. 2013).

There are several important motivations for a model-based semantics of a distributed program. First, we can better separate the program structure, i.e., the rules, from the (distributed) implementation that may change over time. For example, consider rules that generate messages. These rules *can* be implemented with asynchronous communication, but how we evaluate them across machines is eventually just a physical performance decision. Said differently, the point of message rules is not to model a physical phenomenon, but rather to admit a wider array of physical implementations than a local evaluation strategy. Model-based interpretations of a program admit all such implementations, and can perhaps suggest some new ones. Second, we can investigate the *need* for time: we can think about when temporal delay is needed for expressivity, rather than when it is imposed upon us by some implementation detail like physical separation of nodes. In this context we mention the CRON conjecture by Hellerstein, that relates causality on messages to the nature of the computations in which those messages participate (Hellerstein 2010b; Ameloot and den Bussche 2014). We elaborate on causality below.

Concretely, our approach will be to model a distributed program with Datalog under the stable model semantics (Gelfond and Lifschitz 1988) because this semantics is widely used in logic programming. Following the language Dedalus (Alvaro et al. 2009; Alvaro et al. 2011; Hellerstein 2010b), we express the functionality of the distributed program with three kinds of rules: “deductive rules” for local computation, “in-

ductive rules” for persisting memory across local computation steps, and, “asynchronous rules” for representing message sending. The asynchronous rules will non-deterministically choose the arrival times of messages (Krishnamurthy and Naqvi 1988; Saccà and Zaniolo 1990).

However, using only the above rules is not sufficient, as this still allows stable models that express undesirable computations, where messages can be sent “into the past”. Therefore, each program is augmented with a set of rules that express *causality* on the messages. Causality stands for the physical constraint that an effect can only happen after its cause. Applied to message delivery, this intuitively means that a sent message can only be delivered in the future, not in the past. The rules for causality reason from the perspective of the local times of each node, which is a justified approach since there is no common “global clock” in a distributed environment (Attiya and Welch 2004). As a second improvement, we also introduce rules to ensure that only a finite number of messages arrive at each local step of a node, as occurs in a real distributed system. Applying the stable model semantics to the augmented Datalog programs constitutes our modeling of a distributed (Datalog) program.

On another note, it is already well-known that for finite input domains, the combination of Datalog and stable model semantics allows for expressing all problems in NP (Marek and Truszczyński 1999). However, it is not yet clear what can be represented when infinite input domains are considered. From this perspective, our work demonstrates that the stable model semantics is indeed also suitable for modeling distributed programs, whose execution is unbounded in time. Here, time would be provided as an infinite input.

Correctness As we have motivated above, our goal is to describe the workings of a distributed system declaratively, so that new insights can emerge from this perspective. Hence, it is important to verify that the model-based semantics really corresponds to the execution of a distributed program.

To this end, we additionally formalize the execution of a distributed Datalog program by means of an *operational* semantics (Deutsch et al. 2006; Navarro and Rybalchenko 2009; Grumbach and Wang 2010; Ameloot et al. 2011). This second semantics is defined as a transition system. The transition system is infinite because nodes run indefinitely and keep sending messages. In addition, the transition system is highly nondeterministic, because nodes work concurrently and messages can be delayed.

We establish rigorously a correspondence between the features of the operational semantics and the features of the proposed model-based semantics. To formulate our result, we describe each operational execution by a structure that we call a *trace*, which includes for each node in the network the detailed information about the local steps it has performed and about the messages it has sent and received. For our distributed Datalog programs, we show that such operational traces correspond to the set of stable models.

Outline This paper is organized as follows. First, Section 2 discusses related work. Section 3 gives preliminaries. Next, Section 4 represents distributed Datalog pro-

grams under the model-based semantics; this section is based on Dedalus, a Datalog-like language. Section 5 justifies the intuitions of the model-based semantics by establishing an equivalence with an operational semantics. Section 6 finishes with the conclusion.

2 Related Work

The work of Lobo et al. (2012) is closely related to our work. For a Dedalus-inspired language, they give a model-theoretic semantics based on answer set programming, i.e., stable models. To define this semantics, they syntactically translate the rules of their language to Datalog, where all literals are given an explicit location and time variable, to represent the data that each node has during each local time. This translation resembles the model-theoretic semantics for distributed Datalog programs in this paper. To enforce natural execution properties in their semantics, like causality, Lobo et al. specify auxiliary rules in the syntactical translation. The work of Lobo et al. (2012) does not yet mention the connection between the model-theoretic semantics and desired executions of a distributed system, i.e., an operational semantics.

Extending the work of Lobo et al, the work of Ma et al. (2013) formalizes a distributed system as a composition of I/O automata (Lynch 1996). An operational execution of such a system is a sequence of valid transitions, called a trace. Global properties of the system can be analyzed by translating it into a logic program, to which an answer set solver can be applied. Ma et al. mention that operational traces of the system correspond to answer sets of the logic program, and that this provides a formal foundation for the analysis tools based on answer set programming. Thus, the work of Ma et al. (2013) indicates a practical benefit of having a correspondence between a declarative and operational semantics for languages used in declarative networking. As mentioned above, we also establish a similar correspondence in the current paper, for our distributed Datalog programs. We note, however, a few differences between our work and that of Ma et al. First, in the work of Ma et al, the message buffer of a node has a maximum size. In our operational semantics, the buffers are unbounded. Moreover, Ma et al. construct their logic programs for a fixed range of timestamps. In our declarative, model-based semantics, time is given as an infinite input to a Datalog program whose rules are independent of a fixed time range. Lastly, our work devotes much attention to rigorously showing the correspondence between the declarative and operational semantics, whereas this is not elaborated in the work of Ma et al.

Also in the setting of distributed systems, Interlandi et al. (2013) give a Dedalus-inspired language for describing *synchronous* systems. In such systems, the nodes of the network proceed in rounds and the messages can not be arbitrarily delayed. During each round, the nodes share the same global clock. Interlandi et al. specify an operational semantics for their language, based on relational transducer networks (Ameloot et al. 2013). They also show that this operational semantics coincides with a model-theoretic semantics of a single holistic Datalog program. It should be noted that Lobo et al. (2012), and the current paper, deal with *asynchronous*

systems, that in general pose a bigger challenge for a distributed program to be correct, i.e., the program should remain unaffected by nondeterministic effects caused by message delays.

An area of artificial intelligence that is closely related to declarative networking is that of programming multi-agent systems in declarative languages. The knowledge of an agent can be expressed by a logic program, which also allows for non-monotone reasoning, and agents update their knowledge by modifying the rules in these logic programs (Leite et al. 2002; Nigam and Leite 2006; Leite and Soares 2007). The language LUPS (Alferes et al. 2002) was designed to specify such dynamic updates to logic programs, and LUPS is also a declarative language itself. After applying a sequence of updates specified in LUPS, the semantics of the resulting logic program can be defined in an inductive way. But an interesting connection to this current work, is that the semantics can also be given by first syntactically translating the original program and its updates into a single normal logic program, after which the stable model semantics is applied (Alferes et al. 2002). It should be noted however that in this second semantics, there is no modeling of causality or the sending of messages.

Of course, logic programming is not the only means for specifying a (distributed) system. For example, in the area of formal methods, logic-based languages like TLA (Lamport 2000a), Z (Woodcock and Davies 1996), and Event-B (Abrial 2010) can be used to specify various distributed algorithms. Specifications written in these languages can also be automatically checked for correctness.

Although we work within the established setting of declarative networking (Loo et al. 2009), the scientific debate on the merits of Datalog versus other formalisms for programming distributed systems remains open. It seems desirable to have an analysis of how features of Datalog relate to the features of other languages for formal specification, e.g. (Lamport 2000a; Woodcock and Davies 1996; Abrial 2010), both on the syntactical and the semantical level. However, a deep understanding of the other languages would be needed. Moreover, one may expect that features of Datalog will in general not map naturally to features of the other languages. Hence, we consider such a comparison to be a separate research project, outside the scope of the current paper.

3 Preliminaries

3.1 Database Basics

A *database schema* \mathcal{D} is a finite set of pairs (R, k) where R is a *relation name* and $k \in \mathbb{N}$ its associated *arity*. A relation name occurs at most once in a database schema. We often write (R, k) as R/k .

We assume some infinite universe **dom** of atomic data values. A *fact* \mathbf{f} is a pair (R, \bar{a}) , often denoted as $R(\bar{a})$, where R is a relation name and \bar{a} is a tuple of values over **dom**. For a fact $R(\bar{a})$, we call R the *predicate*. We say that a fact $R(a_1, \dots, a_k)$ is *over* database schema \mathcal{D} if $R/k \in \mathcal{D}$. A database *instance* I over \mathcal{D} is a set of facts over \mathcal{D} . For a subset $\mathcal{D}' \subseteq \mathcal{D}$, we write $I|_{\mathcal{D}'}$ to denote the subset of facts in

I whose predicate is a relation name in \mathcal{D}' . We write $adom(I)$ to denote the set of values occurring in facts of I .

3.2 Datalog with Negation

We recall Datalog with negation (Abiteboul et al. 1995), abbreviated Datalog[−]. We assume the standard database perspective, where a Datalog[−] program is evaluated over a given set of facts, i.e., where these facts are not part of the program itself.

Let **var** be a universe of *variables*, disjoint from **dom**. An *atom* is of the form $R(u_1, \dots, u_k)$ where R is a relation name and $u_i \in \mathbf{var} \cup \mathbf{dom}$ for each $i = 1, \dots, k$. We call R the *predicate*. If an atom contains no data values, we call it *constant-free*. A *literal* is an atom or an atom with “ \neg ” prepended. A literal that is an atom is called *positive* and otherwise it is called *negative*.

It will be technically convenient to use a slightly unconventional definition of rules. Formally, a Datalog[−] rule φ is a triple

$$(head_\varphi, pos_\varphi, neg_\varphi)$$

where $head_\varphi$ is an atom; pos_φ and neg_φ are sets of atoms; and, the variables in φ all occur in pos_φ . This last condition is called *safety*. The components $head_\varphi$, pos_φ and neg_φ are called respectively the *head*, the *positive body atoms* and the *negative body atoms*. We refer to $pos_\varphi \cup neg_\varphi$ as the *body atoms*. Note, neg_φ contains just atoms, not negative literals. Every Datalog[−] rule φ must have a head, whereas pos_φ and neg_φ may be empty. If $neg_\varphi = \emptyset$ then φ is called *positive*.

A rule φ may be written in the conventional syntax. For instance, if $head_\varphi = T(u, v)$, $pos_\varphi = \{R(u, v)\}$ and $neg_\varphi = \{S(v)\}$, with $u, v \in \mathbf{var}$, then we can write φ as

$$T(u, v) \leftarrow R(u, v), \neg S(v).$$

The specific ordering of literals to the right of the arrow has no significance in this paper.

The set of variables of φ is denoted $vars(\varphi)$. If $vars(\varphi) = \emptyset$ then φ is called *ground*, in which case $\{head_\varphi\} \cup pos_\varphi \cup neg_\varphi$ is a set of facts.

Let \mathcal{D} be a database schema. A rule φ is said to be *over schema \mathcal{D}* if for each atom $R(u_1, \dots, u_k) \in \{head_\varphi\} \cup pos_\varphi \cup neg_\varphi$ we have $R/k \in \mathcal{D}$. A Datalog[−] program P over \mathcal{D} is a set of (safe) Datalog[−] rules over \mathcal{D} . We write $sch(P)$ to denote the smallest database schema that P is over; note, $sch(P)$ is uniquely defined. We define $idb(P) \subseteq sch(P)$ to be the database schema consisting of all relations in rule-heads of P . We abbreviate $edb(P) = sch(P) \setminus idb(P)$.¹

Any database instance I over $sch(P)$ can be given as *input* to P . Note, I may already contain facts over $idb(P)$.² Let $\varphi \in P$. A *valuation* for φ is a total function $V : vars(\varphi) \rightarrow \mathbf{dom}$. The *application* of V to an atom $R(u_1, \dots, u_k)$ of φ , denoted

¹ The abbreviation “idb” stands for “intensional database schema” and “edb” stands for “extensional database schema” (Abiteboul et al. 1995).

² The need for this will become clear in Section 5.

$V(R(u_1, \dots, u_k))$, results in the fact $R(a_1, \dots, a_k)$ where for each $i \in \{1, \dots, k\}$ we have $a_i = V(u_i)$ if $u_i \in \mathbf{var}$ and $a_i = u_i$ otherwise. In words: applying V replaces the variables by data values and leaves the old data values unchanged. This is naturally extended to a set of atoms, which results in a set of facts. Valuation V is said to be *satisfying for φ on I* if $V(pos_\varphi) \subseteq I$ and $V(neg_\varphi) \cap I = \emptyset$. If so, φ is said to *derive the fact $V(head_\varphi)$* .

3.2.1 Positive and Semi-positive

Let P be a Datalog⁻ program. We say that P is *positive* if all rules of P are positive. We say that P is *semi-positive* if for each rule $\varphi \in P$, the atoms of neg_φ are over $edb(P)$. Note, positive programs are semi-positive.

We now give the semantics of a semi-positive Datalog⁻ program P (Abiteboul et al. 1995). First, let T_P be the *immediate consequence operator* that maps each instance J over $sch(P)$ to the instance $J' = J \cup A$ where A is the set of facts derived by all possible satisfying valuations for the rules of P on J .

Let I be an instance over $sch(P)$. Consider the infinite sequence I_0, I_1, I_2 , etc, inductively defined as follows: $I_0 = I$ and $I_i = T_P(I_{i-1})$ for each $i \geq 1$. The *output of P on input I* , denoted $P(I)$, is defined as $\bigcup_j I_j$; this is the *minimal fixpoint* of the T_P operator. Note, $I \subseteq P(I)$. When I is finite, the fixpoint is finite and can be computed in polynomial time according to data complexity (Vardi 1982).

3.2.2 Stratified Semantics

We now recall the stratified semantics for a Datalog⁻ program P (Abiteboul et al. 1995). As a slight abuse of notation, here we will treat $idb(P)$ as a set of only relation names (without associated arities). First, P is called *syntactically stratifiable* if there is a function $\sigma : idb(P) \rightarrow \{1, \dots, |idb(P)|\}$ such that for each rule $\varphi \in P$, having some head predicate T , the following conditions are satisfied:

- $\sigma(R) \leq \sigma(T)$ for each $R(\bar{u}) \in pos_\varphi|_{idb(P)}$;
- $\sigma(R) < \sigma(T)$ for each $R(\bar{u}) \in neg_\varphi|_{idb(P)}$.

For $R \in idb(P)$, we call $\sigma(R)$ the *stratum number* of R . For technical convenience, we may assume that if there is an $R \in idb(P)$ with $\sigma(R) > 1$ then there is an $S \in idb(P)$ with $\sigma(S) = \sigma(R) - 1$. Intuitively, function σ partitions P into a sequence of semi-positive Datalog⁻ programs P_1, \dots, P_k with $k \leq |idb(P)|$ such that for each $i = 1, \dots, k$, the program P_i contains the rules of P whose head predicate has stratum number i . This sequence is called a *syntactic stratification* of P . We can now apply the *stratified semantics* to P : for an input I over $sch(P)$, we first compute the fixpoint $P_1(I)$, then the fixpoint $P_2(P_1(I))$, etc. The *output of P on input I* , denoted $P(I)$, is defined as $P_k(P_{k-1}(\dots P_1(I) \dots))$. It is well known that the output of P does not depend on the chosen syntactic stratification (if more than one exists). Not all Datalog⁻ programs are syntactically stratifiable.

3.2.3 Stable Model Semantics

We now recall the stable model semantics for a Datalog[−] program P (Gelfond and Lifschitz 1988; Saccà and Zaniolo 1990). Let I be an instance over $sch(P)$. Let $\varphi \in P$. Let V be a valuation for φ whose image is contained in $adom(I) \cup C$, where C is the set of all constants appearing in P . Valuation V does not have to be satisfying for φ on I . Together, V and φ give rise to a ground rule ψ , obtained from φ by replacing each $u \in vars(\varphi)$ with $V(u)$. We call ψ a *ground rule of φ with respect to I* . Let $ground(\varphi, I)$ denote the set of all ground rules of φ with respect to I . The *ground program of P on I* , denoted $ground(P, I)$, is defined as $\bigcup_{\varphi \in P} ground(\varphi, I)$. Note, if $I = \emptyset$, the set $ground(P, I)$ contains only rules whose ground atoms are made with C , or atoms that are nullary.

Let M be another instance over $sch(P)$. We write $ground_M(P, I)$ to denote the program obtained from $ground(P, I)$ as follows:

1. remove every rule $\psi \in ground(P, I)$ for which $neg_\psi \cap M \neq \emptyset$;
2. remove the negative (ground) body atoms from all remaining rules.

Note, $ground_M(P, I)$ is a positive program. We say that M is a *stable model of P on input I* if M is the output of $ground_M(P, I)$ on input I . If so, the semantics of positive Datalog[−] programs implies $I \subseteq M$. Not all Datalog[−] programs have stable models on every input (Gelfond and Lifschitz 1988).

3.3 Network and Distributed Databases

A (computer) *network* is a nonempty finite set \mathcal{N} of *nodes*, which are values in **dom**. Intuitively, \mathcal{N} represents the identifiers of compute nodes involved in a distributed system. Communication channels (edges) are not explicitly represented because we allow a node x to send a message to any node y , as long as x knows about y by means of input relations or received messages. For general distributed or cluster computing, the delivery of messages is handled by the network layer, which is abstracted away. But (Datalog) programs can also describe the network layer itself (Loo et al. 2009; Hellerstein 2010b), in which case we would restrict attention to programs where nodes only send messages to nodes to which they are explicitly linked; these nodes would again be provided as input.

A *distributed database instance* H over a network \mathcal{N} and a database schema \mathcal{D} is a function that maps every node of \mathcal{N} to an ordinary finite database instance over \mathcal{D} . This represents how data over the same schema \mathcal{D} is spread over a network.

As a small example of a distributed database instance, consider the following instance H over a network $\mathcal{N} = \{x, y\}$ and a schema $\mathcal{D} = \{R/1, S/1\}$: $H(x) = \{R(a), S(b)\}$ and $H(y) = \{R(a), S(c)\}$. In words: we put facts $R(a)$ and $S(b)$ at node x , and we put facts $R(a)$ and $S(c)$ at node y . Note that it is possible that the same fact is given to multiple nodes.

4 Model-Based Semantics

Here we describe a class of distributed Datalog[⌊] programs that we give a model-based semantics. First, in Section 4.1, we recall the user language Dedalus, that is based on Datalog[⌊] with annotations, in which the programmer can express the functionality of the distributed program. Next, we discuss how to assign a declarative, model-based semantics to Dedalus programs. This semantics consists of applying the stable model semantics to the Dedalus programs after they are transformed into pure Datalog[⌊] programs, i.e., without annotations. We introduce some auxiliary notations and symbols in Section 4.2. Next, in Section 4.3, we give a basic transformation of Dedalus programs in order to apply the stable model semantics. However, this basic transformation has some shortcomings, that we iteratively correct in Sections 4.4 and 4.5.

4.1 User Language: Dedalus

Our user language for distributed Datalog[⌊] programs is Dedalus (Alvaro et al. 2009; Alvaro et al. 2011; Hellerstein 2010b), here presented as Datalog[⌊] with annotations.³ Essentially, the language represents updatable memory for the nodes of a network and provides a mechanism for communication between these nodes.

4.1.1 Syntax

Let \mathcal{D} be a database schema. We write $\mathbf{B}\{\bar{v}\}$, where \bar{v} is a tuple of variables, to denote any sequence β of literals over database schema \mathcal{D} , such that the variables in β are precisely those in the tuple \bar{v} . Let $R(\bar{u})$ denote any atom over \mathcal{D} . There are three types of Dedalus rules over \mathcal{D} :

- A *deductive* rule is a normal Datalog[⌊] rule over \mathcal{D} .
- An *inductive* rule is of the form

$$R(\bar{u}) \bullet \leftarrow \mathbf{B}\{\bar{u}, \bar{v}\}.$$

- An *asynchronous* rule is of the form

$$R(\bar{u}) \mid y \leftarrow \mathbf{B}\{\bar{u}, \bar{v}, y\}.$$

For asynchronous rules, the annotation ‘ $\mid y$ ’ with $y \in \mathbf{var}$ means that the derived head facts are transferred (“piped”) to the addressee node represented by y . Deductive, inductive and asynchronous rules will express respectively local computation, updatable memory, and message sending. As in Section 3.2, a Dedalus rule is called *safe* if all its variables occur in at least one positive body atom.

We already provide some intuition of how asynchronous rules operate. There are four conceptual time points involved in the execution of an asynchronous rule: the time when the body is evaluated; the time when the derived fact is sent to the addressee; the time when the fact arrives at the addressee; and, the time when

³ These annotations correspond to syntactic sugar in the previous presentations of Dedalus.

the arrived fact becomes visible at the addressee. In the model-based semantics presented later, the first two time points coincide and the last two time points coincide; and, there is no upper bound on the interval between these two pairs, although it will be finite.

Now consider the following definition:

Definition 4.1

A *Dedalus program* over a schema \mathcal{D} is a set of deductive, inductive and asynchronous Dedalus rules over \mathcal{D} , such that all rules are safe, and the set of deductive rules is syntactically stratifiable.

In the current work, we will additionally assume that Dedalus programs are constant-free, as is common in the theory of database query languages, and which is not really a limitation, since constants that are important for the program can always be indicated by unary relations in the input.

Let \mathcal{P} be a Dedalus program. The definitions of $sch(\mathcal{P})$, $idb(\mathcal{P})$, and $edb(\mathcal{P})$ are like for Datalog⁻ programs. An *input* for \mathcal{P} is a *distributed* database instance over some network \mathcal{N} and the schema $edb(\mathcal{P})$.

4.1.2 Semantics Sketch

We sketch the main idea behind the semantics of a Dedalus program \mathcal{P} . We illustrate the semantics in Section 4.1.3.

Let H be an input distributed database instance for \mathcal{P} , over a network \mathcal{N} . The idea is that all nodes $x \in \mathcal{N}$ run the same program \mathcal{P} and use their local input fragment $H(x)$ to do local computation and to send messages. Conceptually, each node of \mathcal{N} should be thought of as doing local computation steps, indefinitely. During each step, a node reads the following facts: (i) the local input; (ii) some received message facts, generated by asynchronous rules on other nodes or the node itself; and, (iii) the facts derived by inductive rules during the previous step on this same node. Next, the deductive rules are applied to these available facts, to compute a fixpoint D under the stratified semantics.

Subsequently, the asynchronous and inductive rules are fired in parallel on the deductive fixpoint D , trying all possible valuations in single-step derivations (i.e., no fixpoint). The asynchronous rules send messages to other nodes or to the same node. Messages arrive after an arbitrary (but finite) delay, where the delay can vary for each message. The inductive rules store facts in the memory of the local node. The effect of an inductive derivation is only visible in the very next step; so, if a fact is to be remembered over multiple steps, it should always be explicitly rederived by inductive rules.

4.1.3 Examples

We consider several examples to demonstrate the three kinds of Dedalus rules, and how they work together. These examples also illustrate the utility of Dedalus when applied to some practical problems. Here, we follow the principle that the output on a node x consists of the facts that are eventually derived during every step of x .

```

marked(u) | y ← start(u), Node(y).

marked(u)• ← marked(u).

marked(v) ← marked(u), R(u, v).
vert(u) ← R(u, v).
vert(u) ← R(v, u).
missing() ← vert(u), ¬marked(u).
covered() ← ¬missing().

```

Figure 1. Dedalus program for Example 1.

Example 1

In this example we compute reachable vertices on graph data. Consider the Dedalus program \mathcal{P} in Figure 1. We assume the *edb* relations $R/2$, **start**/1, and **Node**/1. For each node, relation R describes a local graph, and relation **start** provides certain starting vertices. In any input distributed database instance H over a network \mathcal{N} , we assume that for each node, relation **Node** is initialized to contain all nodes of \mathcal{N} ; intuitively, **Node** can be regarded as an address book for \mathcal{N} .

Now, the idea is that each node of \mathcal{N} will check whether all of its local vertices are reachable from the (distributed) start vertices. Communication is needed to share these start vertices, which is accomplished by the asynchronous rule. The receipt of a start vertex initializes a local relation **marked**/1 at each node; this relation contains reachable vertices. The inductive rule says that all reachable vertices that we know during the current step, are remembered in the next step. This way, the effect of the communication is preserved. Moreover, the third rule, which is deductive, collects all local graph vertices reachable from the currently known reachable vertices. Note, the inductive rule will cause the result of this deductive computation to be also remembered in the next step, although this effect is not really needed here. The last four rules, which are deductive, check that all local vertices are reachable from the start vertices seen so far; if so, a local flag **covered**() is derived.

In our semantics, we will enforce that all messages eventually arrive. In such a semantics, eventually a node will produce **covered**() during each step iff all its local vertices are reachable from the distributed start vertices. \square

Example 2

In this example we generate a random ordering of a set through asynchronous delivery of messages. Every node generates a random ordering of a local *edb* relation $S/1$ that represents an input set. We also assume an *edb* relation **Id**/1 that contains on each node the identifier of that node; the relation **Id** allows a node to send a message to itself. The idea is that a node sends all elements of S to itself as messages, and the arbitrary arrival order is used to generate an ordering of the elements. This

ordering depends on the execution, and some executions will not lead to orderings if some elements are always jointly delivered.

The corresponding program is shown in Figure 2. We use relation $M/1$ to send the elements of S , as accomplished by the single asynchronous rule. The relations $F/1$ and $N/2$ represent the ordering of S so far, and they are considered as the output of the program; the letters ‘F’ and ‘N’ stand for “first” and “next” respectively. For example, a possible ordering of the set $\{a, b, c, d\}$ could be expressed by the following facts: $F(d)$, $N(d, c)$, $N(c, b)$, $N(b, a)$.

Inductive rules are responsible for remembering the iteratively updated versions of F and N . The other rules are deductive, and they can conceptually be executed in the order in which they are written. The main technical challenge is to only update the ordering when precisely one element of S arrives; otherwise, because we have no choice mechanism, we would accidentally give the same ordinal to two different elements. Checking whether we may update the ordering is accomplished through other auxiliary relations. We use a nullary relation `started` as a flag to know whether we still have to initialize relation F or not.

Note that the program keeps sending all elements of S through the single asynchronous rule. Alternatively, by adapting the program, we could send the elements only once by making sure the asynchronous rule is fired only once (in parallel for all elements of S). In that case, as soon as two elements are later delivered together, the ordering will not contain all elements. \square

Example 3

This example is inspired by commit protocols that were expressed in a precursor language of Dedalus (Alvaro et al. 2009). In particular, we implement a two-phase commit protocol where agents, represented by nodes, vote either “yes” or “no” for transaction identifiers. Such a protocol could be part of a bigger system, where transactions are distributed across agents and each agent may only perform the transaction locally if *all* agents want to do this. A single *coordinator* node is responsible for combining the votes for each transaction identifier t : the coordinator broadcasts “yes” for t if all votes for t are “yes”, and “no” otherwise. Each agent stores the decision of the coordinator.

Because the agents and the coordinator have different roles, we make two separate Dedalus programs.⁴ First, the agent nodes are assigned the following simple Dedalus program, whose relations are explained below:

$$\text{vote}(t, x, v) \mid y \leftarrow \text{myVote}(t, v), \text{Id}(x), \text{coord}(y).$$

$$\text{outcome}(t, v) \bullet \leftarrow \text{outcome}(t, v).$$

Here, the *edb* relations are: `myVote/2` that maps each transaction identifier t to a local vote “yes” or “no”, `Id/1` storing the identifier of the agent, and `coord/1`

⁴ In our formal definitions, all nodes execute the same Dedalus program. However, it is easy to simulate two different programs by giving every node the union of both programs, but using a flag to guard the rules of each program. In this example, we can then assume that one node gets a “coordinator” flag as input, and the other nodes get an “agent” flag as input.

```

M(u) | x ← S(u), Id(x).

used(u) ← F(u).
used(u) ← N(u, v).
used(u) ← N(v, u).
new(u) ← M(u), ¬used(u).
eq(u, u) ← S(u).
two() ← new(u), new(v), ¬eq(u, v).
keep(u) ← new(u), ¬two().
notlast(u) ← N(u, v).
last(u) ← F(u), ¬notlast(u).
last(u) ← N(v, u), ¬notlast(u).

started() ← F(u).
F(u)• ← ¬started(), keep(u).
N(u, v)• ← started(), last(u), keep(v).
F(u)• ← F(u).
N(u, v)• ← N(u, v).

```

Figure 2. Dedalus program for Example 2.

storing the identifier of the coordinator. Also, the relations `vote/3` and `outcome/2` represent respectively the outgoing votes and the final decision by the coordinator.

Second, the coordinator node is assigned the Dedalus program shown in Figure 3. The coordinator has the following *edb* relations: relation `T/1` containing all transaction identifiers, relations `Y/1` and `N/1` containing the constants “yes” and “no” respectively, and relation `agents/1` containing all voting agents. The coordinator uses an inductive rule to gradually accumulate all votes for each transaction identifier. Votes can have arbitrary delays, but in our model the delays are always finite. In each computation step, the deductive rules at the coordinator recompute a relation `complete` that contains the transaction identifiers for which all votes have been received. When a transaction identifier t has at least one “no” vote, the coordinator decides “no” for t , and otherwise the coordinator decides “yes” for t . The final decision is broadcast to all agents. The coordinator adds the transactions with a decision to a log, so the decision will not be broadcast again. \square

4.2 Auxiliary Notations and Relations

Let \mathcal{P} be a Dedalus program. Let $R/k \in \text{sch}(\mathcal{P})$. We will use facts of the form $R(x, s, a_1, \dots, a_k)$ to express that fact $R(a_1, \dots, a_k)$ is present at a node x during its local step s , with $s \in \mathbb{N}$, after the deductive rules are executed. We call x the

```

vote(t, x, v)• ← vote(t, x, v).
known(t, x) ← vote(t, x, v).
missing(t) ← T(t), agent(x), ¬known(t, x).
complete(t) ← T(t), ¬missing(t).
decideNo(t) ← votes(t, x, v), N(v).
decideYes(t) ← complete(t), ¬decideNo(t).

outcome(t, v) | y ← decideNo(t), ¬log(t), N(v), agent(y).
outcome(t, v) | y ← decideYes(t), ¬log(t), Y(v), agent(y).
log(t)• ← complete(t).
log(t)• ← log(t).

```

Figure 3. Dedalus (coordinator) program for Example 3.

location specifier and s the *timestamp*. In order to represent timestamps, we assume $\mathbb{N} \subseteq \mathbf{dom}$.

We write $sch(\mathcal{P})^{LT}$ to denote the database schema obtained from $sch(\mathcal{P})$ by incrementing the arity of every relation by two. The two extra components will contain the location specifier and timestamp.⁵ For an instance I over $sch(\mathcal{P})$, $x \in \mathbf{dom}$ and $s \in \mathbb{N}$, we write $I^{\uparrow x, s}$ to denote the facts over $sch(\mathcal{P})^{LT}$ that are obtained by prepending location specifier x and timestamp s to every fact of I . Also, if L is a sequence of literals over $sch(\mathcal{P})$, and $\mathbf{x}, \mathbf{s} \in \mathbf{var}$, we write $L^{\uparrow \mathbf{x}, \mathbf{s}}$ to denote the sequence of literals over $sch(\mathcal{P})^{LT}$ that is obtained by adding location specifier \mathbf{x} and timestamp \mathbf{s} to the literals in L (negative literals stay negative).

We also need auxiliary relation names, that are assumed not to be used in $sch(\mathcal{P})$; these are listed in Table 1.⁶ The concrete purpose of these relations will become clear in the following subsections.

We define the following schema

$$\mathcal{D}_{\text{time}} = \{\text{time}/1, \text{tsucc}/2, </2, \neq/2\}.$$

The relations ‘<’ and ‘≠’ will be written in infix notation in rules. We consider only the following instance over $\mathcal{D}_{\text{time}}$:

$$\begin{aligned}
I_{\text{time}} = & \{\text{time}(s), \text{tsucc}(s, s+1) \mid s \in \mathbb{N}\} \\
& \cup \{(s < t) \mid s, t \in \mathbb{N} : s < t\} \\
& \cup \{(s \neq t) \mid s, t \in \mathbb{N} : s \neq t\}.
\end{aligned}$$

Intuitively, the instance I_{time} provides timestamps together with relations to compare them.

⁵ The abbreviation ‘LT’ stands for “location specifier and timestamp”.

⁶ In practice, auxiliary relations can be differentiated from those in $sch(\mathcal{P})$ by a namespace mechanism.

Table 1. Relation names not in $sch(\mathcal{P})$.

Relation Names	Meaning
all	network
time , tsucc , < , ≠	timestamps
before	happens-before relation
cand_R , chosen_R , other_R , for each relation name R in $idb(\mathcal{P})$	messages
hasSender , isSmaller , hasMax , rcvInf	only a finite number of messages arrive at each step of a node

4.3 Dynamic Choice Transformation

Let \mathcal{P} be a Dedalus program. We describe the *dynamic choice transformation* to transform \mathcal{P} into a pure Datalog[−] program $pure_{ch}(\mathcal{P})$. The most technical part of the transformation involves the use of dynamic choice to select an arrival timestamp for each message generated by an asynchronous rule. The actual transformation is presented first; next we give the semantics; and, lastly, we discuss how the transformation can be improved.

4.3.1 Transformation

We incrementally construct $pure_{ch}(\mathcal{P})$. In particular, for each rule in \mathcal{P} , we specify what corresponding rule (or rules) should be added to $pure_{ch}(\mathcal{P})$. For technical convenience, we assume that rules of \mathcal{P} always contain at least one positive body atom. This assumption allows us to more elegantly enforce that head variables in rules of $pure_{ch}(\mathcal{P})$ also occur in at least one positive body atom.⁷ Let $\mathbf{x}, \mathbf{s}, \mathbf{t}, \mathbf{t}' \in \mathbf{var}$ be distinct variables not yet occurring in rules of \mathcal{P} . We write $\mathbf{B}\{\bar{\mathbf{v}}\}$, where $\bar{\mathbf{v}}$ is a tuple of variables, to denote any sequence β of literals over $sch(\mathcal{P})$, such that the variables in β are precisely those in $\bar{\mathbf{v}}$. Also recall the notations and relation names from Section 4.2.

Deductive rules For each deductive rule $R(\bar{\mathbf{u}}) \leftarrow \mathbf{B}\{\bar{\mathbf{u}}, \bar{\mathbf{v}}\}$ in \mathcal{P} , we add to $pure_{ch}(\mathcal{P})$ the following rule:

$$R(\mathbf{x}, \mathbf{s}, \bar{\mathbf{u}}) \leftarrow \mathbf{B}\{\bar{\mathbf{u}}, \bar{\mathbf{v}}\}^{\uparrow \mathbf{x}, \mathbf{s}}. \quad (1)$$

⁷ This assumption is not really a restriction, since a nullary positive body atom is already sufficient.

This rule expresses that deductively derived facts at some node x during step s are (immediately) visible within step s of x . Note, all atoms in this rule are over $sch(\mathcal{P})^{LT}$.

Inductive rules For each inductive rule $R(\bar{u}) \bullet \leftarrow \mathbf{B}\{\bar{u}, \bar{v}\}$ in \mathcal{P} , we add to $pure_{ch}(\mathcal{P})$ the following rule:

$$R(\mathbf{x}, \mathbf{t}, \bar{u}) \leftarrow \mathbf{B}\{\bar{u}, \bar{v}\}^{\uparrow \mathbf{x}, \mathbf{s}}, \mathbf{tsucc}(\mathbf{s}, \mathbf{t}). \quad (2)$$

This rule expresses that inductively derived facts becomes visible in the *next* step of the *same* node.

Asynchronous rules We use facts of the form $\mathbf{all}(x)$ to say that x is a node of the network at hand. We use facts of the form $\mathbf{cand}_R(x, s, y, t, \bar{a})$ to express that node x at its step s sends a message $R(\bar{a})$ to node y , and that t could be the arrival timestamp of this message at y .⁸ Within this context, we use a fact $\mathbf{chosen}_R(x, s, y, t, \bar{a})$ to say that t is the *effective* arrival timestamp of this message at y . Lastly, a fact $\mathbf{other}_R(x, s, y, t, \bar{a})$ means that t is *not* the arrival timestamp of the message. Now, for each asynchronous rule

$$R(\bar{u}) \mid y \leftarrow \mathbf{B}\{\bar{u}, \bar{v}, y\}$$

in \mathcal{P} , letting \bar{w} be a tuple of new and distinct variables with $|\bar{w}| = |\bar{u}|$, we add to $pure_{ch}(\mathcal{P})$ the following rules, for which the intuition is given below:

$$\mathbf{cand}_R(\mathbf{x}, \mathbf{s}, \mathbf{y}, \mathbf{t}, \bar{u}) \leftarrow \mathbf{B}\{\bar{u}, \bar{v}, y\}^{\uparrow \mathbf{x}, \mathbf{s}}, \mathbf{all}(\mathbf{y}), \mathbf{time}(\mathbf{t}). \quad (3)$$

$$\mathbf{chosen}_R(\mathbf{x}, \mathbf{s}, \mathbf{y}, \mathbf{t}, \bar{w}) \leftarrow \mathbf{cand}_R(\mathbf{x}, \mathbf{s}, \mathbf{y}, \mathbf{t}, \bar{w}), \neg \mathbf{other}_R(\mathbf{x}, \mathbf{s}, \mathbf{y}, \mathbf{t}, \bar{w}). \quad (4)$$

$$\mathbf{other}_R(\mathbf{x}, \mathbf{s}, \mathbf{y}, \mathbf{t}, \bar{w}) \leftarrow \mathbf{cand}_R(\mathbf{x}, \mathbf{s}, \mathbf{y}, \mathbf{t}, \bar{w}), \mathbf{chosen}_R(\mathbf{x}, \mathbf{s}, \mathbf{y}, \mathbf{t}', \bar{w}), \mathbf{t} \neq \mathbf{t}'. \quad (5)$$

$$R(\mathbf{y}, \mathbf{t}, \bar{w}) \leftarrow \mathbf{chosen}_R(\mathbf{x}, \mathbf{s}, \mathbf{y}, \mathbf{t}, \bar{w}). \quad (6)$$

Rule (3) represents the messages that are sent. It evaluates the body of the original asynchronous rule, verifies that the addressee is within the network by using relation \mathbf{all} , and it generates all possible candidate arrival timestamps.

Now remains the matter of actually choosing *one* arrival timestamp amongst all these candidates. Intuitively, rule (4) selects an arrival timestamp for a message with the condition that this timestamp is not yet ignored, as expressed with relation \mathbf{other}_R . Also, looking at rule (5), a possible arrival timestamp t becomes ignored if there is already a chosen arrival timestamp t' with $t \neq t'$. Together, both rules have the effect that exactly one arrival timestamp will be chosen under the stable model semantics. This technical construction is due to Saccà and Zaniolo (1990), who show how to express dynamic choice under the stable model semantics.

Rule (6) represents the actual arrival of an R -message with the chosen arrival timestamp: the data-tuple in the message becomes part of the addressee's state for relation R . When the addressee reads relation R , it thus transparently reads the arrived R -messages.

⁸ Here, ' \mathbf{cand} ' abbreviates "candidate".

Note, if multiple asynchronous rules in \mathcal{P} have the same head predicate R , only new \mathbf{cand}_R -rules have to be added because the rules (4)–(6) are general for all R -messages.

Note that if there are asynchronous rules in \mathcal{P} , program $\mathit{pure}_{\text{ch}}(\mathcal{P})$ is not syntactically stratifiable if a \mathbf{cand}_R -rule contains a body atom that (indirectly) negatively depends on R .⁹ In that case, $\mathit{pure}_{\text{ch}}(\mathcal{P})$ might not even be locally stratifiable (Apt and Bol 1994).

4.3.2 Semantics

Now we define the semantics of $\mathit{pure}_{\text{ch}}(\mathcal{P})$. Let H be an input distributed database instance for \mathcal{P} , over a network \mathcal{N} . Using the notations from Section 4.2, we define $\mathit{decl}(H)$ to be the following database instance over the schema $\mathit{edb}(\mathcal{P})^{\text{LT}} \cup \{\mathbf{all}/1\} \cup \mathcal{D}_{\text{time}}$:

$$\begin{aligned} \mathit{decl}(H) = & \{R(x, s, \bar{a}) \mid x \in \mathcal{N}, s \in \mathbb{N}, R(\bar{a}) \in H(x)\} \\ & \cup \{\mathbf{all}(x) \mid x \in \mathcal{N}\} \cup I_{\text{time}}. \end{aligned}$$

In words: we make for each node its input facts available at all timestamps; we provide the set of all nodes; and, I_{time} provides the timestamps with comparison relations.¹⁰ Note, instance $\mathit{decl}(H)$ is infinite because \mathbb{N} is infinite.

The stable model semantics for Datalog[−] programs is reviewed in Section 3.2.3. Consider now the following definition:

Definition 4.2

For an input distributed database instance H for \mathcal{P} , we call any stable model of $\mathit{pure}_{\text{ch}}(\mathcal{P})$ on input $\mathit{decl}(H)$ a *choice-model* of \mathcal{P} on input H .

4.3.3 Possible Improvement

We illustrate a shortcoming of the dynamic choice transformation. Consider the Dedalus program \mathcal{P} in Figure 4. We assume that in each input distributed database, the edb relation $\mathbf{Id}/1$ contains on each node just the identifier of this node. This way, the node can send messages to itself. Relation T is the intended output relation of \mathcal{P} . The idea is that a node sends $A()$ to itself continuously. When $A()$ arrives, we send $B()$, but we also want to create an output fact $T()$. We only create $T()$ when $B()$ is absent. When $B()$ is received, it is remembered by inductive rules. Now, we see that the delivery of at least one $A()$ is necessary to cause a $B()$ to be sent. This creates the expectation that $T()$ is always created: at least one $A()$ is delivered before any $B()$. This intuition can be formalized as *causality* (Attiya and Welch 2004) (see also Section 5.2.1).

However, this intuition is violated by some choice-models of \mathcal{P} , as we demonstrate

⁹ Indeed, \mathbf{cand}_R is used to compute R , but R is also used to compute \mathbf{cand}_R , giving a cycle through negation.

¹⁰ For simplicity we already include relation $<$ in this definition, although this relation will only be used later.

$ \begin{aligned} &A() \mid \mathbf{x} \leftarrow \text{Id}(\mathbf{x}). \\ &B() \mid \mathbf{x} \leftarrow A(), \text{Id}(\mathbf{x}). \\ &T() \leftarrow A(), \neg B(). \\ &T() \bullet \leftarrow T(). \\ &B() \bullet \leftarrow B(). \end{aligned} $

Figure 4. Dedalus program sensitive to non-causality.

next. Consider the input distributed database instance H over a singleton network $\{z\}$ that assigns the fact $\text{Id}(z)$ to z . Now, consider the following choice-model M of \mathcal{P} on H :¹¹

$$M = \text{decl}(H) \cup M_A^{\text{snd}} \cup M_A^{\text{rcv}} \cup M_B^{\text{snd}} \cup M_B^{\text{rcv}},$$

where

$$\begin{aligned}
M_A^{\text{snd}} = & \{ \text{cand}_A(z, s, z, t) \mid s, t \in \mathbb{N} \} \\
& \cup \{ \text{chosen}_A(z, s, z, s+1) \mid s \in \mathbb{N} \} \\
& \cup \{ \text{other}_A(z, s, z, t) \mid s, t \in \mathbb{N}, t \neq s+1 \};
\end{aligned}$$

$$M_A^{\text{rcv}} = \{ A(z, s) \mid s \in \mathbb{N}, s \geq 1 \};$$

$$\begin{aligned}
M_B^{\text{snd}} = & \{ \text{cand}_B(z, s, z, t) \mid s, t \in \mathbb{N}, s \geq 1 \} \\
& \cup \{ \text{chosen}_B(z, 1, z, 0) \} \\
& \cup \{ \text{chosen}_B(z, s, z, s+1) \mid s \in \mathbb{N}, s \geq 2 \} \\
& \cup \{ \text{other}_B(z, 1, z, t) \mid t \in \mathbb{N}, t \neq 0 \} \\
& \cup \{ \text{other}_B(z, s, z, t) \mid s, t \in \mathbb{N}, s \geq 2, t \neq s+1 \};
\end{aligned}$$

$$M_B^{\text{rcv}} = \{ B(z, s) \mid s \in \mathbb{N} \}.$$

In M_B^{snd} , note that one B -message is sent at timestamp 1 of z , and arrives at timestamp 0 of z . We immediately see that this message is peculiar: we should not be able to send a message to arrive in the past. Because of the stray message $B()$, the fact $B()$ exists at all timestamps: it arrives at timestamp 0 and is henceforth persisted by the inductive rule for relation B ; this is modeled by set M_B^{rcv} . Subsequently, there are no ground rules of the form $T(z, s) \leftarrow A(z, s)$ with $s \in \mathbb{N}$ in the ground program $\text{ground}_M(C, I)$, where $C = \text{pure}_{\text{ch}}(\mathcal{P})$ and $I = \text{decl}(H)$.

In the next subsection, we exclude such unintuitive stable models using an extended transformation of Dedalus programs.

¹¹ Using straightforward arguments, it can indeed be shown that M is a stable model of $\text{pure}_{\text{ch}}(\mathcal{P})$ on $\text{decl}(H)$.

4.4 Causality Transformation

Let \mathcal{P} be a Dedalus program. In this section, we present the causality transformation $pure_{ca}(\mathcal{P})$ that extends $pure_{ch}(\mathcal{P})$ to exclude the unintuitive stable models that we have encountered in the previous subsection. We first present the new transformation, and then we discuss how the transformation can still be improved.

4.4.1 Transformation

We define $pure_{ca}(\mathcal{P})$ again incrementally. First, we transform deductive and inductive rules just as in $pure_{ch}(\mathcal{P})$.

Next, we use facts of the form **before**(x, s, y, t) to express that local step s of node x happens before local step t of node y . Regardless of \mathcal{P} , we always add the following rules to $pure_{ca}(\mathcal{P})$:

$$\mathbf{before}(x, s, x, t) \leftarrow \mathbf{all}(x), \mathbf{tsucc}(s, t) \cdot \quad (7)$$

$$\mathbf{before}(x, s, y, t) \leftarrow \mathbf{before}(x, s, z, u), \mathbf{before}(z, u, y, t) \cdot \quad (8)$$

Rule (7) expresses that on every node, a step happens before the next step. Rule (8) makes relation **before** transitive.

Now, for each asynchronous rule

$$R(\bar{u}) \mid y \leftarrow \mathbf{B}\{\bar{u}, \bar{v}, y\}$$

in \mathcal{P} , we add to $pure_{ca}(\mathcal{P})$ the previous transformation rules (4), (5) and (6) (omitting the **cand_R**-rule), and we add the following new rules, where \bar{w} is a tuple of new and distinct variables with $|\bar{w}| = |\bar{u}|$, and x, s , and t are also new variables:

$$\mathbf{cand}_R(x, s, y, t, \bar{u}) \leftarrow \mathbf{B}\{\bar{u}, \bar{v}, y\}^{\uparrow x, s}, \mathbf{all}(y), \mathbf{time}(t), \neg \mathbf{before}(y, t, x, s) \cdot \quad (9)$$

$$\mathbf{before}(x, s, y, t) \leftarrow \mathbf{chosen}_R(x, s, y, t, \bar{w}) \cdot \quad (10)$$

Like the old rule (3), rule (9) represents the messages that are sent, but now candidate arrival timestamps are restricted by relation **before** to enforce causality. Intuitively, this restriction prevents cycles from occurring in relation **before**. This aligns with the semantics of a real distributed system, where the happens-before relation is a strict partial order (Attiya and Welch 2004) (see also Section 5.2.1).

Rule (10) adds the causal restriction that the local step of the sender happens before the arrival step of the addressee. Together with the previously introduced rules (7) and (8), this will make sure that when the addressee later *causally* replies to the sender, the reply — as generated by a rule of the form (9) — will arrive after this first send-step of the sender.

Remark 1

The new program $pure_{ca}(\mathcal{P})$ excludes unintuitive models like the one in Section 4.3.3. In the context of that particular example, it will be impossible to exhibit a stable model of $pure_{ca}(\mathcal{P})$ in which $B()$ is sent to timestamp 0. Indeed, $B()$ can only be

sent starting from timestamp 1; timestamp 0 at z (locally) happens before timestamp 1 at z ; and, the negative **before**-literal in rule (9) will prevent sending from timestamp 1 at z to timestamp 0 at z . Also in scenarios where different nodes x and y send messages to each other, when node x replies to a message of node y sent at timestamp s of y , node x can not send the reply to a timestamp t of y with $t < s$.

4.4.2 Semantics

The semantics of the causality transformation is the same as for the dynamic choice transformation:

Definition 4.3

For an input distributed database instance H for \mathcal{P} , we call any stable model of $\text{pure}_{\text{ca}}(\mathcal{P})$ on input $\text{decl}(H)$ a *causal model* of \mathcal{P} on input H .

4.4.3 Possible Improvement

We illustrate a shortcoming of the causality transformation. Consider the Dedalus program \mathcal{P} in Figure 5. We assume that in each input distributed database, the *edb* relation **contact**/1 contains intended recipients of messages. Relation T serves as the output relation of \mathcal{P} . The idea is that a node sends $A()$ to its recipients continuously. When $A()$ arrives, a recipient sets a local flag **first**(). Later, when a second $A()$ arrives, the recipient creates an output fact $T()$ that we remember by means of inductive rules. Intuitively, we expect that $T()$ is always created because the fact $A()$ is sent infinitely often to a recipient, making this recipient witness the arrival of $A()$ at (hopefully) two distinct moments.

However, this intuition is violated by some causal models of \mathcal{P} . Consider the input distributed database instance H over a network $\{x, y\}$ that (only) assigns the fact **contact**(y) to x . Now, consider the following causal model M of \mathcal{P} on H :¹²

$$M = \text{decl}(H) \cup M_A^{\text{snd}} \cup M_A^{\text{rcv}} \cup M^{\text{before}},$$

where

$$\begin{aligned} M_A^{\text{snd}} = & \quad \{\text{cand}_A(x, s, y, t) \mid s, t \in \mathbb{N}\} \\ & \cup \{\text{chosen}_A(x, s, y, 0) \mid s \in \mathbb{N}\} \\ & \cup \{\text{other}_A(x, s, y, t) \mid s, t \in \mathbb{N}, t \neq 0\}; \end{aligned}$$

$$\begin{aligned} M_A^{\text{rcv}} = & \quad \{A(y, 0)\} \\ & \cup \{\text{first}(y, s) \mid s \in \mathbb{N}, s \geq 1\}; \end{aligned}$$

$$\begin{aligned} M^{\text{before}} = & \quad \{\text{before}(x, s, x, t) \mid s, t \in \mathbb{N}, s < t\}; \\ & \cup \{\text{before}(y, s, y, t) \mid s, t \in \mathbb{N}, s < t\}; \\ & \cup \{\text{before}(x, s, y, t) \mid s, t \in \mathbb{N}\} \end{aligned}$$

¹² Using straightforward arguments, it can be shown that M is a stable model of $\text{pure}_{\text{ca}}(\mathcal{P})$ on $\text{decl}(H)$.

```

A() | y ← contact(y).

first()• ← A().
first()• ← first().

T() ← first(), A().
T()• ← T().

```

Figure 5. Dedalus program sensitive to infinite message grouping.

In this causal model, all instances of message $A()$ that x sends to y arrive at timestamp 0 of y . For this reason, node y can not witness two different arrivals of message $A()$. In practice, however, node y can not receive an *infinite* number of messages during a timestamp, and the deliveries of the $A()$ messages would be spread out more evenly in time. So, in the next subsection, we will additionally exclude such infinite message arrivals, to obtain our final transformation of Dedalus programs.

4.5 Causality-Finiteness Transformation

Let \mathcal{P} be a Dedalus program. As seen in the previous subsection, program $pure_{ca}(\mathcal{P})$ allows an infinite number of messages to arrive at any step of a node. This does not happen in any real-world distributed system; indeed, no node has to process an infinite number of messages at any given moment. We consider this to be an additional restriction that must be explicitly enforced. To this purpose, we present in this section the causality-finiteness transformation $pure(\mathcal{P})$ that extends $pure_{ca}(\mathcal{P})$.

We will approach this problem as follows. Suppose there are an infinite number of messages that arrive at some node y during its step t . Since in a network there are only a finite number of nodes and a node can only send a finite number of messages during each step (the input domain is finite), there must be at least one node x that sends messages to step t of y during an infinite number of steps of x . Hence there is no maximum value amongst the corresponding send-timestamps of x . Thus, in order to prevent the arrival of an infinite number of messages at step t of y , it will be sufficient to demand that there always is such a maximum send-timestamp for every sender. Below, we will implement this strategy with some concrete rules in $pure(\mathcal{P})$.

4.5.1 Transformation

We define $pure(\mathcal{P})$ as $pure_{ca}(\mathcal{P})$ extended as follows. The additional rules can be thought of as being relative to an addressee and a step of this addressee, represented by the variables y and t respectively.

We use a fact $rcvInf(y, t)$ to express that node y receives an infinite number of messages during its step t . First, we add the following rule to $pure(\mathcal{P})$ for *each*

relation chosen_R that results from the transformation of asynchronous rules in $\text{pure}_{\text{ca}}(\mathcal{P})$, where \mathbf{x} , \mathbf{s} , \mathbf{y} , and \mathbf{t} are variables and $\bar{\mathbf{w}}$ is a tuple of distinct variables disjoint from the previous ones with $|\bar{\mathbf{w}}|$ the arity of relation R in $\text{sch}(\mathcal{P})$:

$$\text{hasSender}(\mathbf{y}, \mathbf{t}, \mathbf{x}, \mathbf{s}) \leftarrow \text{chosen}_R(\mathbf{x}, \mathbf{s}, \mathbf{y}, \mathbf{t}, \bar{\mathbf{w}}), \neg \text{rcvInf}(\mathbf{y}, \mathbf{t}). \quad (11)$$

This rule intuitively means that as long as addressee \mathbf{y} has not received an infinite number of messages during its step \mathbf{t} , we register the senders and their send-timestamps.

Recall the auxiliary relations defined in Section 4.2. Next, we add to $\text{pure}(\mathcal{P})$ the following rules, for which the intuition is provided below:

$$\text{isSmaller}(\mathbf{y}, \mathbf{t}, \mathbf{x}, \mathbf{s}) \leftarrow \text{hasSender}(\mathbf{y}, \mathbf{t}, \mathbf{x}, \mathbf{s}), \text{hasSender}(\mathbf{y}, \mathbf{t}, \mathbf{x}, \mathbf{s}'), \quad (12)$$

$$\mathbf{s} < \mathbf{s}'.$$

$$\text{hasMax}(\mathbf{y}, \mathbf{t}, \mathbf{x}) \leftarrow \text{hasSender}(\mathbf{y}, \mathbf{t}, \mathbf{x}, \mathbf{s}), \neg \text{isSmaller}(\mathbf{y}, \mathbf{t}, \mathbf{x}, \mathbf{s}). \quad (13)$$

$$\text{rcvInf}(\mathbf{y}, \mathbf{t}) \leftarrow \text{hasSender}(\mathbf{y}, \mathbf{t}, \mathbf{x}, \mathbf{s}), \neg \text{hasMax}(\mathbf{y}, \mathbf{t}, \mathbf{x}). \quad (14)$$

Rule (12) checks for each sender and each of its send-timestamps whether there is a later send-timestamp of that same sender. Rule (13) tries to find a maximum send-timestamp. Finally, rule (14) derives a rcvInf -fact if no maximum send-timestamp was found for at least one sender.

We will show in Section 5.3.1 that in any stable model, the above rules make sure that every node receives only a finite number of messages at every step.

4.5.2 Semantics

The semantics of the causality-finiteness transformation is again the same as for the dynamic choice transformation and the causality transformation:

Definition 4.4

For an input distributed database instance H for \mathcal{P} , we call any stable model of $\text{pure}(\mathcal{P})$ on input $\text{decl}(H)$ a *causal-finite model* of \mathcal{P} on input H .

We will refer to a causal-finite model also simply as *model*.

5 Correctness

In Section 4, we have described the computation of a distributed Datalog⁺ program by means of stable models. By using suitable rules, we have excluded some unintuitive stable models. But at this point we are still not sure whether the remaining stable models really correspond to the execution of a distributed system. We fill that gap in this section: we show that each remaining stable model corresponds to an execution of the distributed Datalog⁺ program under an operational semantics, and vice versa. We call such an execution a *run*, and we will only be concerned with so-called *fair* runs, where each node is made active infinitely often and all sent messages are eventually delivered.

We extract from each run \mathcal{R} a *trace*, denoted $\text{trace}(\mathcal{R})$, which is a set of facts

that shows in detail what each node computes during each step. We will make this concrete in the following subsections. But we can already state our main result, as follows:

Theorem 4

Let \mathcal{P} be a Dedalus program. For each input distributed database instance H for \mathcal{P} ,

- (i) for every fair run \mathcal{R} of \mathcal{P} there is a model M of \mathcal{P} such that $\text{trace}(\mathcal{R}) = M|_{\text{sch}(\mathcal{P})^{\text{LT}}}$, and
- (ii) for every model M of \mathcal{P} there is a fair run \mathcal{R} of \mathcal{P} such that $\text{trace}(\mathcal{R}) = M|_{\text{sch}(\mathcal{P})^{\text{LT}}}$.

□

First, Section 5.1 formalizes runs and traces of runs. The proof of item (i) of the theorem is described in Section 5.2. The proof of item (ii), which is the most difficult, is described in Section 5.3. We only describe the crucial reasoning steps of the proofs; the intricate technical details can be found in the online appendix to the paper.

5.1 Operational Semantics

In this section, we give an operational semantics for Dedalus that is in line with earlier formal work on declarative networking (Deutsch et al. 2006; Navarro and Rybalchenko 2009; Grumbach and Wang 2010; Ameloot et al. 2011; Abiteboul et al. 2011).

Let \mathcal{P} be a Dedalus program, and let H be an input distributed database instance for \mathcal{P} , over a network \mathcal{N} . The essence of the operational semantics is as follows. Every node of \mathcal{N} runs program \mathcal{P} , and a node has access only to its own local state and any received messages. The nodes are made active one by one in some arbitrary order, and this continues an infinite number of times. During each active moment of a node x , called a *local (computation) step*, node x receives message facts and applies its deductive, inductive and asynchronous rules. Concretely, the deductive rules, forming a stratified Datalog[−] subprogram, are applied to the incoming messages and the previous state of x . Next, the inductive rules are applied to the output of the deductive subprogram, and these allow x to store facts in its memory: these facts become visible in the next local step of x . Finally, the asynchronous rules are also applied to the output of the deductive subprogram, and these allow x to send facts to the other nodes or to itself. These facts become visible at the addressee after some arbitrary delay, which represents asynchronous communication, as occurs for instance on the Internet. We assume that all messages are eventually delivered (and are thus never lost). We will refer to local steps simply as “steps”.

We make the above sketch more concrete in the next subsections.

5.1.1 Configurations

Let \mathcal{P} , H , and \mathcal{N} be as above. A configuration describes the network at a certain point in its evolution. Formally, a *configuration* of \mathcal{P} on H is a pair $\rho = (st, bf)$ where

- st is a function mapping each node of \mathcal{N} to an instance over $sch(\mathcal{P})$; and,
- bf is a function mapping each node of \mathcal{N} to a set of pairs of the form (i, \mathbf{f}) , where $i \in \mathbb{N}$ and \mathbf{f} is a fact over $idb(\mathcal{P})$.

We call st and bf the *state* and (*message*) *buffer* respectively. The state says for each node what facts it has stored in its memory, and the message buffer bf says for each node what messages have been sent to it but that are not yet received. The reason for having numbers i , called *send-tags*, attached to facts in the image of bf is merely a technical convenience: these numbers help separate multiple instances of the same fact when it is sent at different moments (to the same addressee), and these send-tags will not be visible to the Dedalus program. For example, if the buffer of a node x simultaneously contains pairs $(3, \mathbf{f})$ and $(7, \mathbf{f})$, this means that \mathbf{f} was sent to x during the operational network transitions with indices 3 and 7, and that both particular instances of \mathbf{f} are not yet delivered to x . This will become more concrete in Section 5.1.3.

The *start configuration* of \mathcal{P} on input H , denoted $start(\mathcal{P}, H)$, is the configuration $\rho = (st, bf)$ defined by $st(x) = H(x)$ and $bf(x) = \emptyset$ for each $x \in \mathcal{N}$. In words: for every node, the state is initialized with its local input fragment in H , and there are no sent messages.

5.1.2 Subprograms

We look at the operations that are executed locally during each step of a node. We have mentioned that the three types of Dedalus rules each have their own purpose in the operational semantics. For this reason, we split the program \mathcal{P} into three subprograms, that contain respectively the deductive, inductive and asynchronous rules. In Section 5.1.3, we describe how these subprograms are used in the operational semantics.

- First, we define $deduc_{\mathcal{P}}$ to be the Datalog[−] program consisting of precisely all deductive rules of \mathcal{P} .
- Secondly, we define $induc_{\mathcal{P}}$ to be the Datalog[−] program consisting of all inductive rules of \mathcal{P} after the annotation ‘•’ in their head is removed.
- Thirdly, we define $async_{\mathcal{P}}$ to be the Datalog[−] program consisting of precisely all rules

$$T(\mathbf{y}, \bar{\mathbf{u}}) \leftarrow \mathbf{B}\{\bar{\mathbf{u}}, \mathbf{y}\}$$

where

$$T(\bar{\mathbf{u}}) \mid \mathbf{y} \leftarrow \mathbf{B}\{\bar{\mathbf{u}}, \mathbf{y}\}$$

is an asynchronous rule of \mathcal{P} . So, we basically put the variable \mathbf{y} as the first

component in the (extended) head atom. The intuition for the generated head facts is that the first component will represent the addressee.

Note that the programs $deduc_{\mathcal{P}}$, $induc_{\mathcal{P}}$ and $async_{\mathcal{P}}$ are just Datalog⁺ programs over the schema $sch(\mathcal{P})$, or a subschema thereof. Moreover, $deduc_{\mathcal{P}}$ is syntactically stratifiable because the deductive rules in every Dedalus program must be syntactically stratifiable. It is possible however that $induc_{\mathcal{P}}$ and $async_{\mathcal{P}}$ are not syntactically stratifiable. Now we define the semantics of each of these three subprograms.

Let I be a database instance over $sch(\mathcal{P})$. During each step of a node, the intuition of the deductive rules is that they “complete” the available facts by adding all new facts that can be logically derived from them. This calls for a fixpoint semantics, and for this reason, we define the *output of $deduc_{\mathcal{P}}$ on input I* , denoted as $deduc_{\mathcal{P}}(I)$, to be given by the stratified semantics. This implies $I \subseteq deduc_{\mathcal{P}}(I)$. Importantly, I is allowed to contain facts over $idb(\mathcal{P})$, and the intuition is that these facts were derived during a previous step (by inductive rules) or received as messages (as sent by asynchronous rules). This will become more explicit in Section 5.1.3.

During each step of a node, the intuition behind the inductive rules is that they store facts in the memory of the node, and these stored facts will become visible during the next step. There is no notion of a fixpoint here because facts that will become visible in the next step are not available in the current step to derive more facts. For this reason, we define the *output of $induc_{\mathcal{P}}$ on input I* to be the set of facts derived by the rules of $induc_{\mathcal{P}}$ for all possible satisfying valuations in I , in just one derivation step. This output is denoted as $induc_{\mathcal{P}}(I)$.

During each step of a node, the intuition behind the asynchronous rules is that they generate message facts that are to be sent around the network. The *output for $async_{\mathcal{P}}$ on input I* is defined in the same way as for $induc_{\mathcal{P}}$, except that we now use the rules of $async_{\mathcal{P}}$ instead of $induc_{\mathcal{P}}$. This output is denoted as $async_{\mathcal{P}}(I)$. The intuition for not requiring a fixpoint for $async_{\mathcal{P}}$ is that a message fact will arrive at another node, or at a later step of the sender node, and can therefore not be read during sending.

Regarding data complexity (Vardi 1982), for each subprogram the output can be computed in PTIME with respect to the size of its input.

5.1.3 Transitions and Runs

Transitions formalize how to go from one configuration to another. Here we use the subprograms of \mathcal{P} . Transitions are chained to form a *run*. Regarding notation, for a set m of pairs of the form (i, \mathbf{f}) , we define $untag(m) = \{\mathbf{f} \mid \exists i \in \mathbb{N} : (i, \mathbf{f}) \in m\}$.

A *transition with send-tag $i \in \mathbb{N}$* is a five-tuple $(\rho_a, x, m, i, \rho_b)$ such that $\rho_a = (st_a, bf_a)$ and $\rho_b = (st_b, bf_b)$ are configurations of \mathcal{P} on input H , $x \in \mathcal{N}$, $m \subseteq bf_a(x)$, and, letting

$$\begin{aligned} I &= st_a(x) \cup untag(m), \\ D &= deduc_{\mathcal{P}}(I), \\ \delta^{i \rightarrow y} &= \{(i, R(\bar{a})) \mid R(y, \bar{a}) \in async_{\mathcal{P}}(D)\} \text{ for each } y \in \mathcal{N}, \end{aligned}$$

for x and each $y \in \mathcal{N} \setminus \{x\}$ we have

$$\begin{aligned} st_b(x) &= H(x) \cup \text{induc}_{\mathcal{P}}(D), & st_b(y) &= st_a(y), \\ bf_b(x) &= (bf_a(x) \setminus m) \cup \delta^{i \rightarrow x}, & bf_b(y) &= bf_a(y) \cup \delta^{i \rightarrow y}. \end{aligned}$$

We call ρ_a and ρ_b respectively the *source* and *target* configuration, and say this transition is *of* the *active* node x . Intuitively, the transition expresses that x reads its old state together with the received facts in $\text{untag}(m)$ (thus without the tags), and describes the subsequent computation: subprogram $\text{deduc}_{\mathcal{P}}$ completes the available information; the new state of x consists of the input facts of x united with all facts derived by subprogram $\text{induc}_{\mathcal{P}}$; and, subprogram $\text{async}_{\mathcal{P}}$ generates *messages*, whose first component indicates the addressee.¹³ Note, $\text{induc}_{\mathcal{P}}$ and $\text{async}_{\mathcal{P}}$ do not influence each other, and can be thought of as being executed in parallel. Also, for each $y \in \mathcal{N}$, the set $\delta^{i \rightarrow y}$ contains all messages addressed to y , with send-tag i attached. Messages with an addressee outside the network are ignored. This way of defining local computation closely corresponds to that of the language Webdamlog (Abiteboul et al. 2011). If $m = \emptyset$, we call the transition a *heartbeat*.

A *run* \mathcal{R} of \mathcal{P} on input H is an infinite sequence of transitions, such that (i) the source configuration of the first transition is $\text{start}(\mathcal{P}, H)$, (ii) the target configuration of each transition is the source configuration of the next transition, and (iii) the transition at ordinal i of the sequence uses send-tag i . Ordinals start at 0 for technical convenience. The resulting transition system is highly non-deterministic because in each transition we can choose the active node and also what messages to deliver; the latter choice is represented by the set m from above.

Remark 2 (Parallel transitions)

Transitions as defined here can simulate *parallel* transitions in which multiple nodes are active at the same time and receive messages from their respective buffers. Indeed, if we would have multiple nodes active during a parallel transition, they would receive messages from their buffers in isolation, and this can be represented by a chain of transitions in which these nodes receive one after the other precisely the messages that they received in the parallel transition. For this reason, we limit our attention to transitions with single active nodes.

5.1.4 Fairness and Arrival Function

In the literature on process models it is customary to require certain fairness conditions on the execution of a system, for instance to exclude some extreme situations that are expected not to happen in reality (Francez 1986; Apt et al. 1988; Lamport 2000b).

Let \mathcal{R} be a run of \mathcal{P} on H . For every transition $i \in \mathbb{N}$, let $\rho_i = (st_i, bf_i)$ denote the source configuration of transition i . Now, \mathcal{R} is called *fair* if:

- every node is the active node in an infinite number of transitions of \mathcal{R} ; and,

¹³ Note, input facts are preserved by the transition. This aligns with the design of Dedalus, where we do not allow facts to be retracted; only negation as failure is permitted.

- for every transition $i \in \mathbb{N}$, for every $y \in \mathcal{N}$, for every pair $(j, \mathbf{f}) \in bf_i(y)$, there is a transition k with $i \leq k$ in which (j, \mathbf{f}) is delivered to y .

Intuitively, the fairness conditions disallow starvation: every node does an infinite number of local computation steps and every sent message is eventually delivered. We consider only fair runs in this paper. Note, a fair run exists for every input because heartbeats remain possible even when there are no messages to deliver.

In the second condition about message deliveries, it is possible that $k = i$, and in that case (j, \mathbf{f}) is delivered in the transition immediately following configuration ρ_i . Because the pair (j, \mathbf{f}) can be in the message buffer of multiple nodes, this k is not unique for the pair (j, \mathbf{f}) by itself. But, when we also consider the addressee y , it follows from the operational semantics that this k is unique for the triple (j, y, \mathbf{f}) .

This reasoning gives rise to a function $\alpha_{\mathcal{R}}$, called the *arrival function for \mathcal{R}* , that is defined as follows: for every transition i , for every node y , for every message \mathbf{f} sent to addressee y during i , the function $\alpha_{\mathcal{R}}$ maps (i, y, \mathbf{f}) to the transition ordinal k in which (i, \mathbf{f}) is delivered to y . We always have $\alpha_{\mathcal{R}}(i, y, \mathbf{f}) > i$. Indeed, the delivery of a message can only happen after it was sent. So, when the delivery of one message causes another to be sent, then the second one is delivered in a later transition. This is related to the topic of causality that we have introduced in Section 4. This topic will also be further discussed in Sections 5.2 and 5.3.

5.1.5 Timestamps and Trace

For each transition i of a run, we define the *timestamp* of the active node x during i to be the number of transitions of x that come strictly before i . This can be thought of as the *local* (zero-based) clock of x during i , and is denoted $loc_{\mathcal{R}}(i)$. For example, suppose we have the following sequence of active nodes: x, y, y, x, x , etc. If we would write the timestamps next to the nodes, we get this sequence: $(x, 0), (y, 0), (y, 1), (x, 1), (x, 2)$, etc.

As a counterpart to function $loc_{\mathcal{R}}(\cdot)$, for each $(x, s) \in \mathcal{N} \times \mathbb{N}$ we define $glob_{\mathcal{R}}(x, s)$ to be the transition ordinal i of \mathcal{R} such that x is the active node in transition i and $loc_{\mathcal{R}}(i) = s$. In words: we find the transition in which node x does its local computation step with timestamp s . It follows from the definition of $loc_{\mathcal{R}}(\cdot)$ that $glob_{\mathcal{R}}(x, s)$ is uniquely defined.

Let \mathcal{R} be a run of \mathcal{P} on input H . Recall that H is over network \mathcal{N} . We now capture the computed data during \mathcal{R} as a set of facts that we call the *trace*. For each transition $i \in \mathbb{N}$, let x_i denote the active node, and let D_i denote the output of subprogram $deduc_{\mathcal{P}}$ during i . The operational semantics implies that D_i consists of (i) the input *edb*-facts at x_i ; (ii) the inductively derived facts during the previous step of x_i (if $loc_{\mathcal{R}}(i) \geq 1$); (iii) the messages delivered during transition i ; and, (iv) all facts deductively derived from the previous ones. So, intuitively, D_i contains *all* local facts over $sch(\mathcal{P})$ that x_i has during transition i .

Recall the notations of Section 4.2. Now, the trace of \mathcal{R} is the following instance over $sch(\mathcal{P})^{LT}$:

$$trace(\mathcal{R}) = \bigcup_{i \in \mathbb{N}} D_i^{\uparrow x_i, loc_{\mathcal{R}}(i)}.$$

The trace shows in detail what happens in the run, in terms of what facts are available on the nodes during which of their steps.

5.2 Run to Model

Let \mathcal{P} be a Dedalus program and let H be an input distributed database instance for \mathcal{P} , over a network \mathcal{N} . Let \mathcal{R} be a fair run of \mathcal{P} on input H . We show there is a model M of \mathcal{P} on H such that $trace(\mathcal{R}) = M|_{sch(\mathcal{P})^{LT}}$. The main idea is that we translate the transitions of \mathcal{R} to facts over the schema of $pure(\mathcal{P})$.

First, in Section 5.2.1, we extract the happens-before relation on nodes and timestamps from \mathcal{R} . Next, in Section 5.2.2, we define the desired model M .

5.2.1 Happens-before Relation

In the operational semantics, we order the actions of the nodes on a fine-grained global time axis, by ordering the transitions in the runs. By contrast, we now define a partial order on $\mathcal{N} \times \mathbb{N}$, saying which steps of nodes must have come before which steps of (other) nodes, without referring to the global ordering imposed by transitions.

First, we extract from \mathcal{R} the message sending and receiving events. Formally, we define $mesg(\mathcal{R})$ to be the set of all tuples (x, s, y, t, \mathbf{f}) , with \mathbf{f} a fact, and denoting $i = glob_{\mathcal{R}}(x, s)$ and $j = glob_{\mathcal{R}}(y, t)$, such that $\alpha_{\mathcal{R}}(i, y, \mathbf{f}) = j$, i.e., node x during step s sends message \mathbf{f} to y that arrives at the step t of y , with possibly $x = y$. In words: $mesg(\mathcal{R})$ contains the direct relationships between local steps of nodes that arise through message sending.

From \mathcal{R} we can now extract the *happens-before* relation (Attiya and Welch 2004) on the set $\mathcal{N} \times \mathbb{N}$, which is defined as the smallest relation $\prec_{\mathcal{R}}$ on $\mathcal{N} \times \mathbb{N}$ that satisfies the following three conditions:

- for each $(x, s) \in \mathcal{N} \times \mathbb{N}$, we have $(x, s) \prec_{\mathcal{R}} (x, s + 1)$;
- $(x, s) \prec_{\mathcal{R}} (y, t)$ whenever for some fact \mathbf{f} we have $(x, s, y, t, \mathbf{f}) \in mesg(\mathcal{R})$;
- $\prec_{\mathcal{R}}$ is transitive, i.e., $(x, s) \prec_{\mathcal{R}} (z, u) \prec_{\mathcal{R}} (y, t)$ implies $(x, s) \prec_{\mathcal{R}} (y, t)$.

We call these three cases respectively *local* edges, *message* edges and *transitive* edges. Naturally, the first two cases express a direct relationship, whereas the third case is more indirect.

Note, if two runs on the same input have the same happens-before relation, they do not necessarily have the same trace. This is because relation $\prec_{\mathcal{R}}$ does not talk about the specific messages that arrive at the nodes.

We will now show that $\prec_{\mathcal{R}}$ is a strict partial order. Consider first the following property:

Lemma 1

For every run \mathcal{R} , for each $(x, s) \in \mathcal{N} \times \mathbb{N}$ and $(y, t) \in \mathcal{N} \times \mathbb{N}$, if $(x, s) \prec_{\mathcal{R}} (y, t)$ then $glob_{\mathcal{R}}(x, s) < glob_{\mathcal{R}}(y, t)$.

Proof

We can consider a path from (x, s) to (y, t) in $\prec_{\mathcal{R}}$. We can substitute each transitive edge in this path with a subpath of non-transitive edges. This results in a path of only non-transitive edges:

$$(x_1, s_1) \prec_{\mathcal{R}} (x_2, s_2) \prec_{\mathcal{R}} \dots \prec_{\mathcal{R}} (x_n, s_n),$$

where $n \geq 2$, $(x_1, s_1) = (x, s)$ and $(x_n, s_n) = (y, t)$. Because there are no transitive edges, for each $i \in \{1, \dots, n-1\}$, the edge $(x_i, s_i) \prec_{\mathcal{R}} (x_{i+1}, s_{i+1})$ falls into one of the following two cases:

- $x_i = x_{i+1}$ and $s_{i+1} = s_i + 1$ (local edge);
- x_i during step s_i sends a message to x_{i+1} that arrives in step s_{i+1} of x_{i+1} (message edge).

In the first case, it follows from the definition of $loc_{\mathcal{R}}(\cdot)$ that

$$glob_{\mathcal{R}}(x_i, s_i) < glob_{\mathcal{R}}(x_{i+1}, s_{i+1}).$$

For the second case, by our operational semantics, every message is always delivered in a later transition than the one in which it was sent. So, again we have

$$glob_{\mathcal{R}}(x_i, s_i) < glob_{\mathcal{R}}(x_{i+1}, s_{i+1}).$$

Since this property holds for all the above edges, by transitivity we thus have $glob_{\mathcal{R}}(x, s) < glob_{\mathcal{R}}(y, t)$, as desired. \square

Corollary 1

For every run \mathcal{R} , the relation $\prec_{\mathcal{R}}$ is a strict partial order on $\mathcal{N} \times \mathbb{N}$.

Proof

From its definition, we immediately have that $\prec_{\mathcal{R}}$ is transitive. Secondly, irreflexivity for $\prec_{\mathcal{R}}$ follows from Lemma 1. \square

5.2.2 Definition of M

Now we define the model M :

$$M = decl(H) \cup \bigcup_{i \in \mathbb{N}} trans_{\mathcal{R}}^{[i]},$$

where $trans_{\mathcal{R}}^{[i]}$ for each $i \in \mathbb{N}$ is an instance over the schema of $pure(\mathcal{P})$ that describes transition i of \mathcal{R} .¹⁴ Let $i \in \mathbb{N}$. We define $trans_{\mathcal{R}}^{[i]}$ as

$$trans_{\mathcal{R}}^{[i]} = caus_{\mathcal{R}}^{[i]} \cup fin_{\mathcal{R}}^{[i]} \cup duc_{\mathcal{R}}^{[i]} \cup snd_{\mathcal{R}}^{[i]},$$

¹⁴ Note, M must include the input $decl(H)$ by definition of stable model (see Section 3.2.3).

where each of these sets focuses on different aspects of transition i , and they are defined next. Regarding notation, let $\prec_{\mathcal{R}}$ be the happens-before relation as defined in the preceding subsection; let $loc_{\mathcal{R}}(\cdot)$, $glob_{\mathcal{R}}(\cdot)$, and $\alpha_{\mathcal{R}}$ be as defined in Section 5.1; let x_i denote the active node of transition i ; and, let us abbreviate $s_i = loc_{\mathcal{R}}(i)$.

Causality We define $caus_{\mathcal{R}}^{[i]}$ to consist of all facts **before** (x, s, x_i, s_i) for which $(x, s) \in \mathcal{N} \times \mathbb{N}$ and $(x, s) \prec_{\mathcal{R}} (x_i, s_i)$. Intuitively, $caus_{\mathcal{R}}^{[i]}$ represents the joint result of rules (7), (8), and (10), corresponding to respectively the local edges, transitive edges, and message edges of $\prec_{\mathcal{R}}$.

Finite Messages We define $fin_{\mathcal{R}}^{[i]}$ to represent that only a finite number of messages are delivered in transition i , thus at step s_i of node x_i . We proceed as follows. First, let $senders_{\mathcal{R}}^{[i]}$ be the set of all pairs $(x, s) \in \mathcal{N} \times \mathbb{N}$ such that, denoting $j = glob_{\mathcal{R}}(x, s)$, for some fact \mathbf{f} we have $\alpha_{\mathcal{R}}(j, x_i, \mathbf{f}) = i$, i.e., the node x during its step s sends a message to x_i with arrival timestamp s_i . It follows from the operational semantics that for each $(x, s) \in senders_{\mathcal{R}}^{[i]}$ we have $glob_{\mathcal{R}}(x, s) < i$. Now, we define $fin_{\mathcal{R}}^{[i]}$ to consist of the following facts:

- the fact **hasSender** (x_i, s_i, x, s) for each $(x, s) \in senders_{\mathcal{R}}^{[i]}$, representing the result of rule (11);
- the fact **isSmaller** (x_i, s_i, x, s) for each $(x, s) \in senders_{\mathcal{R}}^{[i]}$ and $(x, s') \in senders_{\mathcal{R}}^{[i]}$ with $s < s'$, representing the result of rule (12); and,
- the fact **hasMax** (x_i, s_i, x) for each sender-node x mentioned in $senders_{\mathcal{R}}^{[i]}$, representing the result of rule (13).

We know that in \mathcal{R} only a finite number of messages arrive at step s_i of x_i . Hence, we add no fact **rcvInf** (x_i, s_i) to $fin_{\mathcal{R}}^{[i]}$. This also explains why the specification of the **hasMax**-facts above is relatively simple: there is always a maximum send-timestamp for each sender-node.

Deductive Let D_i denote the output of subprogram *deducP* during transition i . We define $duc_{\mathcal{R}}^{[i]}$ to consist of the facts $D_i^{\uparrow x_i, s_i}$. Intuitively, $duc_{\mathcal{R}}^{[i]}$ represents all facts over $sch(\mathcal{P})$ that are available at x_i during step s_i , i.e., the joint result of rules in $pure(\mathcal{P})$ of the form (1), (2) and (6).

Sending We define $snd_{\mathcal{R}}^{[i]}$ to represent the sending of messages during transition i . We proceed as follows. Let $msg_{\mathcal{R}}^{[i]}$ denote the output of subprogram *asyncP* during transition i , restricted to the facts having their addressee-component in the network. Now, we define $snd_{\mathcal{R}}^{[i]}$ to consist of the following facts:

- all facts **cand** $_R(x_i, s_i, y, t, \bar{a})$ for which $R(y, \bar{a}) \in msg_{\mathcal{R}}^{[i]}$ and $t \in \mathbb{N}$ such that $(y, t) \not\prec_{\mathcal{R}} (x_i, s_i)$, representing the result of rule (9);
- all facts **chosen** $_R(x_i, s_i, y, t, \bar{a})$ for which $R(y, \bar{a}) \in msg_{\mathcal{R}}^{[i]}$ and $t = loc_{\mathcal{R}}(j)$ with $j = \alpha_{\mathcal{R}}(i, y, R(\bar{a}))$, representing the result of rule (4); and,
- all facts **other** $_R(x_i, s_i, y, u, \bar{a})$ for which $R(y, \bar{a}) \in msg_{\mathcal{R}}^{[i]}$, $u \in \mathbb{N}$, $(y, u) \not\prec_{\mathcal{R}} (x_i, s_i)$ and $u \neq loc_{\mathcal{R}}(j)$ with $j = \alpha_{\mathcal{R}}(i, y, R(\bar{a}))$, representing the result of rule (5).

Conclusion We can show that M is indeed a model of \mathcal{P} on input H ; this proof can be found in Appendix A of the online appendix to the paper. By construction of M , we have, as desired:

$$M|_{sch(\mathcal{P})^{LT}} = \bigcup_{i \in \mathbb{N}} \text{duc}_{\mathcal{R}}^{[i]} = \bigcup_{i \in \mathbb{N}} D_i^{\uparrow x_i, s_i} = \text{trace}(\mathcal{R}).$$

5.3 Model to Run

Let \mathcal{P} be a Dedalus program and let H be an input distributed database instance for \mathcal{P} , over some network \mathcal{N} . Let M be a model of \mathcal{P} on input H . We show there is a fair run \mathcal{R} of \mathcal{P} on input H such that $\text{trace}(\mathcal{R}) = M|_{sch(\mathcal{P})^{LT}}$.

The direction shown in Section 5.2 is perhaps the most intuitive direction because we only have to show that a concrete set of facts is actually a stable model. In this section we do not yet understand what M can contain. So, a first important step is to show that M has some desirable properties which allow us to construct a run from it.

Using the notation from Section 3.2.3, let G abbreviate the ground program $\text{ground}_M(C, I)$ where $C = \text{pure}(\mathcal{P})$ and $I = \text{decl}(H)$. By definition of M as a stable model, we have $M = G(I)$.

First, it is important to know that in M we find location specifiers where we expect location specifiers and we find timestamps where we expect timestamps. Formally, we call M *well-formed* if:

- for each $R(x, s, \bar{a}) \in M|_{sch(\mathcal{P})^{LT}}$ we have $x \in \mathcal{N}$ and $s \in \mathbb{N}$;
- for each $\text{before}(x, s, y, t) \in M$, we have $x, y \in \mathcal{N}$ and $s, t \in \mathbb{N}$;
- for each fact $\text{cand}_R(x, s, y, t, \bar{a})$, $\text{chosen}_R(x, s, y, t, \bar{a})$ and $\text{other}_R(x, s, y, t, \bar{a})$ in M , we have $x, y \in \mathcal{N}$ and $s, t \in \mathbb{N}$;
- for each fact $\text{hasSender}(x, s, y, t)$, $\text{isSmaller}(x, s, y, t)$, $\text{hasMax}(x, s, y)$ and $\text{rcvInf}(x, s)$ in M , we have $x, y \in \mathcal{N}$ and $s, t \in \mathbb{N}$.

It can be shown by induction on the fixpoint computation of G that M is always well-formed. We omit the details.

The rest of this subsection is organized as follows. In Section 5.3.1, we extract a happens-before relation \prec_M from M . Next, in Section 5.3.2, we construct a run \mathcal{R} : we use \prec_M to establish a total order on $\mathcal{N} \times \mathbb{N}$ that tells us which are the active nodes in the transitions of \mathcal{R} . Finally, we show in Section 5.3.3 that \mathcal{R} is fair.

5.3.1 Partial Order

We define the following relation \prec_M on $\mathcal{N} \times \mathbb{N}$: for each $(x, s) \in \mathcal{N} \times \mathbb{N}$ and $(y, t) \in \mathcal{N} \times \mathbb{N}$, we write $(x, s) \prec_M (y, t)$ if and only if $\text{before}(x, s, y, t) \in M$. The rest of this section is dedicated to showing that \prec_M is a well-founded strict partial order on $\mathcal{N} \times \mathbb{N}$.

Let G abbreviate the ground program $\text{ground}_M(C, I)$ where $C = \text{pure}(\mathcal{P})$ and $I = \text{decl}(H)$. Regarding terminology, an edge $(x, s) \prec_M (y, t)$ is called a *local edge*, a *message edge* or a *transitive edge* if the fact $\text{before}(x, s, y, t) \in M$ can be derived

by a ground rule in G of respectively the form (7), the form (10), or the form (8).¹⁵ It is possible that an edge is of two or even three types at the same time.

Consider the following claim:

Claim 5

Relation \prec_M is a strict partial order on $\mathcal{N} \times \mathbb{N}$.

Proof

We show that \prec_M is transitive and irreflexive.

Transitive First, we show that \prec_M is transitive. Suppose we have $(x, s) \prec_M (z, u)$ and $(z, u) \prec_M (y, t)$. We have to show that $(x, s) \prec_M (y, t)$. By definition of \prec_M , we have $\mathbf{before}(x, s, z, u) \in M$ and $\mathbf{before}(z, u, y, t) \in M$. Because rule (8) is positive, we have the following ground rule in G :

$$\mathbf{before}(x, s, y, t) \leftarrow \mathbf{before}(x, s, z, u), \mathbf{before}(z, u, y, t).$$

Because M is a stable model and the body of the previous ground rule is in M , we obtain $\mathbf{before}(x, s, y, t) \in M$. Hence, $(x, s) \prec_M (y, t)$, as desired.

Irreflexive Because an edge $(x, s) \prec_M (x, s)$ for any $(x, s) \in \mathcal{N} \times \mathbb{N}$ would form a cycle of length one, it is sufficient to show that there are no cycles in \prec_M at all. This gives us irreflexivity, as desired.

First, let \prec'_M denote the restriction of \prec_M to the edges that are local or message edges. Note that this definition allows some edges in \prec'_M to also be transitive. The edges that are missing from \prec'_M with respect to \prec_M are only derivable by ground rules of the form (8); we call these the *pure* transitive edges. We start by showing that \prec'_M contains no cycles. We show this with a proof by contradiction. So, suppose that there is a cycle in $\mathcal{N} \times \mathbb{N}$ through the edges of \prec'_M :

$$(x_1, s_1) \prec_M (x_2, s_2) \prec_M \dots \prec_M (x_n, s_n)$$

with $n \geq 2$ and $(x_1, s_1) = (x_n, s_n)$. We have $\mathbf{before}(x_i, s_i, x_{i+1}, s_{i+1}) \in M$ for each $i \in \{1, \dots, n-1\}$. Based on these **before**-facts, ground rules in G of the form (8) will have derived $\mathbf{before}(x_i, s_i, x_j, s_j) \in M$ for each $i, j \in \{1, \dots, n\}$.

If each edge on the above cycle would be only local, then for each $i, j \in \{1, \dots, n\}$ with $i < j$ we have $x_i = x_j$ and $s_i < s_j$, and hence $s_1 \neq s_n$, which is false. So, there has to be some $k \in \{1, \dots, n-1\}$ such that $(x_k, s_k) \prec_M (x_{k+1}, s_{k+1})$ is a message edge, derived by a ground rule of the form (10):

$$\mathbf{before}(x_k, s_k, x_{k+1}, s_{k+1}) \leftarrow \mathbf{chosen}_R(x_k, s_k, x_{k+1}, s_{k+1}, \bar{a}).$$

Therefore $\mathbf{chosen}_R(x_k, s_k, x_{k+1}, s_{k+1}, \bar{a}) \in M$. This **chosen**_R-fact must be derived by a ground rule of the form (4) in G , which implies that

$$\mathbf{cand}_R(x_k, s_k, x_{k+1}, s_{k+1}, \bar{a}) \in M.$$

This **cand**_R-fact must in turn be derived by a ground rule ψ of the form (9).

¹⁵ The body of such a ground rule has to be in M .

Because rules of the form (9) in $\text{pure}(\mathcal{P})$ contain a negative **before**-atom in their body, the presence of ψ in G requires that $\text{before}(x_{k+1}, s_{k+1}, x_k, s_k) \notin M$. But that is a contradiction, because $\text{before}(x_i, s_i, x_j, s_j) \in M$ for each $i, j \in \{1, \dots, n\}$ (see above).

Now we show there are no cycles in the entire relation \prec_M . Since $M = G(\text{decl}(H))$, we have $M = \bigcup_{i \in \mathbb{N}} M_i$ where $M_0 = \text{decl}(H)$ and $M_i = T(M_{i-1})$ for each $i \geq 1$ where T is the immediate consequence operator of G . By induction on i , we show that an edge $\text{before}(x, s, y, t) \in M_i$ either is a local or message edge, or it can be replaced by a path of local or message edges in M_i . Then any cycle in \prec_M would imply there is a cycle in \prec'_M , which is impossible. So, \prec_M can not contain cycles. Now, this induction property is satisfied for the base case because M_0 does not contain **before**-facts. For the induction hypothesis, assume the property holds for M_{i-1} , where $i \geq 1$. For the inductive step, let $\text{before}(x, s, y, t) \in M_i \setminus M_{i-1}$. If this fact is derived by a ground rule of the form (7) or (10) then the property is satisfied. Now suppose the fact is derived by a ground rule of the form (8):

$$\text{before}(x, s, y, t) \leftarrow \text{before}(x, s, z, u), \text{before}(z, u, y, t).$$

Both body facts are in M_{i-1} , implying M_{i-1} contains a path of local or message edges from (x, s) to (z, u) and from (z, u) to (y, t) . Hence, using $M_{i-1} \subseteq M_i$, the edge $\text{before}(x, s, y, t) \in M_i$ can be replaced by a path of local or message edges in M_i . \square

In Section 4.5 we have added extra rules to $\text{pure}(\mathcal{P})$ to enforce that every node only receives a finite number of messages during each step. We now verify that this works correctly:

Claim 6

For each $(y, t) \in \mathcal{N} \times \mathbb{N}$ there are only a finite number of pairs $(x, s) \in \mathcal{N} \times \mathbb{N}$ such that $(x, s) \prec_M (y, t)$ is a message edge.

Proof

We start by noting that M does not contain the fact $\text{rcvInf}(y, t)$. Indeed, in order to derive this fact, we need a ground rule in G of the form (14), which has a body fact of the form $\text{hasSender}(y, t, x, s)$. Such **hasSender**-facts must be generated by ground rules in G of the form (11). The rule (11) negatively depends on relation rcvInf . Thus, specifically, if we want a ground rule in G that can derive $\text{hasSender}(y, t, x, s)$, we should require the absence of $\text{rcvInf}(y, t)$ from M . So $\text{rcvInf}(y, t) \in M$ requires $\text{rcvInf}(y, t) \notin M$, which is impossible.

The rest of the proof works towards a contradiction. So, suppose that (y, t) has an infinite number of incoming message edges. Because there are only a finite number of nodes in \mathcal{N} , there has to be a node x that has an infinite number of timestamps s such that $\text{before}(x, s, y, t) \in M$ is a message edge. Since it is a message edge, such a fact $\text{before}(x, s, y, t)$ can be generated by a ground rule in G of the form (10), which implies that there is a relation R in $\text{idb}(\mathcal{P})$ and a tuple \bar{a} such that $\text{chosen}_R(x, s, y, t, \bar{a}) \in M$. Because $\text{rcvInf}(y, t) \notin M$ (see above), for each of

these chosen_R -facts, there is a ground rule of the form (11) in M that derives $\text{hasSender}(y, t, x, s) \in M$.

Rule (14) has a negative hasMax -atom in its body. If we can show that $\text{hasMax}(y, t, x) \notin M$, then there will be a ground rule in G of the form (14), where $\text{hasSender}(y, t, x, s) \in M$:

$$\text{rcvInf}(y, t) \leftarrow \text{hasSender}(y, t, x, s).$$

This then causes $\text{rcvInf}(y, t) \in M$, giving the desired contradiction.

Also towards a proof by contradiction, suppose that $\text{hasMax}(y, t, x) \in M$. This means that there is a ground rule ψ in G of the form (13):

$$\text{hasMax}(y, t, x) \leftarrow \text{hasSender}(y, t, x, s).$$

Because the rule (13) contains a negative isSmaller -atom in the body, and because $\psi \in G$, we know that $\text{isSmaller}(y, t, x, s) \notin M$. But because there are infinitely many facts of the form $\text{hasSender}(y, t, x, s') \in M$, there is at least one fact $\text{hasSender}(y, t, x, s') \in M$ with $s < s'$. Moreover, the rule (12) is positive, and therefore the following ground rule is always in G :

$$\text{isSmaller}(y, t, x, s) \leftarrow \text{hasSender}(y, t, x, s), \text{hasSender}(y, t, x, s'), s < s'.$$

Since the body of this ground rule is in M , the rule derives $\text{isSmaller}(y, t, x, s) \in M$, which gives the desired contradiction. \square

An ordering \prec on a set A is called *well-founded* if for each $a \in A$, there are only a finite number of elements $b \in A$ such that $b \prec a$. We now use Claim 6 to show:

Claim 7

Relation \prec_M on $\mathcal{N} \times \mathbb{N}$ is well-founded.

Proof

Let $(x, s) \in \mathcal{N} \times \mathbb{N}$. We have to show that there are only a finite number of pairs $(y, t) \in \mathcal{N} \times \mathbb{N}$ such that $(y, t) \prec_M (x, s)$. Technically, we can limit our attention to paths in \prec_M consisting of local edges and message edges, because if we can show that there are only a finite number of predecessors of (x, s) on such paths, then there are only a finite number of predecessors when we include the transitive edges as well. First we show that every pair $(y, t) \in \mathcal{N} \times \mathbb{N}$ has only a finite number of incoming local and message edges. If $t > 0$, we can immediately see that (y, t) has precisely one incoming local edge, as created by a ground rule of the form (7), and if $t = 0$ then (y, t) has no incoming local edge. Also, Claim 6 tells us that (y, t) has only a finite number of incoming message edges. So, the number of incoming local and message edges in (y, t) is finite.

Let $(y, t) \in \mathcal{N} \times \mathbb{N}$ be a pair such that $(y, t) \prec_M (x, s)$ is a local edge or a message edge. Starting in (x, s) , we can follow this edge backwards so that we reach (y, t) . If (y, t) itself has incoming local or message edges, from (y, t) we can again follow an edge backwards. This way we can incrementally construct backward paths starting from (x, s) . Because at each pair of $\mathcal{N} \times \mathbb{N}$ there are only a finite number of incoming local or message edges (shown above), if (x, s) would have an infinite number of

predecessors, we must be able to construct a backward path of infinite length. We now show that the existence of such an infinite path leads to a contradiction. So, suppose that there is a backward path of infinite length. Because there are only a finite number of nodes in the network \mathcal{N} , there must be a node y that occurs infinitely often on this path. We will now show that, as we progress further along the backward path, we must see the local timestamps of y strictly decrease. Hence, we must eventually reach timestamp 0 of y , after which we cannot decrement the timestamps of y anymore, and thus it is impossible that y occurs infinitely often along the path. Suppose that the timestamps of y do not strictly decrease. There are two cases. First, if the same pair (y, t) would occur twice on the path, we would have a cycle in \prec_M , which is not possible by Claim 5. Secondly, suppose that there are two timestamps t and t' of y such that $t < t'$ and (y, t) occurs before (y, t') on the backward path, meaning that (y, t) lies closer to (x, s) . Because the edges were followed in reverse, we have

$$(y, t') \prec_M \dots \prec_M (y, t).$$

But since $t < t'$, by means of local edges, we always have

$$(y, t) \prec_M (y, t+1) \prec_M \dots \prec_M (y, t').$$

So, there would be a cycle between (y, t') and (y, t) . But that is again impossible by Claim 5. \square

5.3.2 Construction of Run

Let \prec_M be the well-founded strict partial order on $\mathcal{N} \times \mathbb{N}$ as defined in the preceding subsection. The relation \prec_M has the intuition of a happens-before relation of a run (Section 5.2.1), but the novelty is that it comes from a purely declarative model M . We will now use \prec_M to construct a run \mathcal{R} such that $\text{trace}(\mathcal{R}) = M|_{\text{sch}(\mathcal{P})^{\text{LT}}}$.

Total order It is well-known that a well-founded strict partial order can be extended to a well-founded strict total order. So, let $<_M$ be a well-founded strict total order on $\mathcal{N} \times \mathbb{N}$ that extends \prec_M , i.e., for each $(x, s) \in \mathcal{N} \times \mathbb{N}$ and $(y, t) \in \mathcal{N} \times \mathbb{N}$, if $(x, s) \prec_M (y, t)$ then $(x, s) <_M (y, t)$, but the reverse does not have to hold.

Ordering the set $\mathcal{N} \times \mathbb{N}$ according to $<_M$ gives us a sequence of pairs that will form the transitions in the constructed run \mathcal{R} . Concretely, we obtain a sequence of nodes by taking the node-component from each pair. This will form our sequence of active nodes. Similarly, by taking the timestamp-component from each pair of $\mathcal{N} \times \mathbb{N}$, we obtain a sequence of timestamps. These are the local clocks of the active nodes during their transitions.

We introduce some extra notations to help us reason about the ordering of time that is implied by $<_M$. For each $(x, s) \in \mathcal{N} \times \mathbb{N}$, let $\text{glob}_M(x, s) \in \mathbb{N}$ denote the ordinal of (x, s) as implied by $<_M$, which is well-defined because $<_M$ is well-founded. For technical convenience, we let ordinals start at 0. Note, $\text{glob}_M(\cdot)$ is an injective function. For any $i \in \mathbb{N}$, we define (x_i, s_i) to be the unique pair in $\mathcal{N} \times \mathbb{N}$ such that $\text{glob}_M(x_i, s_i) = i$.

As a counterpart to function $glob_M(\cdot)$, for each $i \in \mathbb{N}$ and each $x \in \mathcal{N}$, let $loc_M(i, x)$ denote the size of the set

$$\{s \in \mathbb{N} \mid glob_M(x, s) < i\}.$$

Intuitively, if i is regarded to be the ordinal of a transition in a run, $loc_M(i, x)$ is the number of local steps of x that came before transition i , i.e., the number of transitions before i in which x was the active node. If $x = x_i$ (the active node) then $loc_M(i, x)$ is effectively the timestamp of x during transition i , and if $x \neq x_i$ then $loc_M(i, x)$ is the next timestamp of x that still has to come after transition i . Note, the functions $glob_M(\cdot)$ and $loc_M(\cdot)$ closely resemble the functions $glob_{\mathcal{R}}(\cdot)$ and $loc_{\mathcal{R}}(\cdot)$ of Section 5.1.5.

Configurations We will now define the desired run \mathcal{R} of \mathcal{P} on H . First we define an infinite sequence of configurations ρ_0, ρ_1, ρ_2 , etc. In a second step we will connect each pair of subsequent configurations by a transition. Recall from Section 5.1.1 that a configuration describes for each node what facts it has stored locally (state), and also what messages have been sent to this node but that are not yet received (message buffer). The facts that are stored on a node are either input *edb*-facts, or facts derived by inductive rules in a previous step of the node. The first kind of facts can be easily obtained from M by keeping only the facts over schema $edb(\mathcal{P})^{LT}$, which gives a subset of $decl(H)$.

For the second kind of state facts, we look at the inductively derived facts in M . Rules in $pure(\mathcal{P})$ that represent inductive rules of \mathcal{P} are recognizable as rules of the form (2): they have a head atom over $sch(\mathcal{P})^{LT}$ and they have a (positive) *tsucc*-atom in their body. No other kind of rule in $pure(\mathcal{P})$ has this form. Hence, the ground rules in G that are based on rules of the form (2) are also easily recognizable, and we will call these *inductive ground rules*. A ground rule $\psi \in G$ is called *active* on M if $pos_\psi \subseteq M$, which implies $head_\psi \in M$ because M is stable. Let M^{ind} denote all head atoms of inductive ground rules in G that are active on M . Note that $M^{ind} \subseteq M$. Regarding notation, for an instance I over $sch(\mathcal{P})^{LT}$, we write I^\Downarrow to denote the set $\{R(\bar{a}) \mid \exists x, s : R(x, s, \bar{a}) \in I\}$, and we write $I|^{x,s}$ to denote the set $\{R(y, t, \bar{a}) \in I \mid y = x, t = s\}$.

Now, for each $i \in \mathbb{N}$, for each node $x \in \mathcal{N}$, denoting $s = loc_M(i, x)$, in configuration $\rho_i = (st_i, bf_i)$, the state $st_i(x)$ is defined as

$$((M|_{edb(\mathcal{P})^{LT}})|^{x,s} \cup M^{ind}|^{x,s})^\Downarrow.$$

We remove the location specifier and timestamp because we have to obtain facts over the schema of \mathcal{P} , not over the schema of $pure(\mathcal{P})$.

Now we define the message buffers in the configurations. Recall that the message buffer of a node always contains pairs of the form (j, \mathbf{f}) , where $j \in \mathbb{N}$ is the transition in which fact \mathbf{f} was sent. For each $i \in \mathbb{N}$, for each node $x \in \mathcal{N}$, in configuration $\rho_i = (st_i, bf_i)$, the message buffer $bf_i(x)$ is defined as

$$\{(glob_M(y, t), R(\bar{a})) \mid \exists u : \mathbf{chosen}_R(y, t, x, u, \bar{a}) \in M, \\ glob_M(y, t) < i \leq glob_M(x, u)\}.$$

Note the use of addressee x in this definition. The definition of $bf_i(x)$ reflects the operational semantics, in that the messages in the buffer of node x must be sent in a previous transition, as expressed by the constraint $glob_M(y, t) < i$. Moreover, the constraint $i \leq glob_M(x, u)$ says that $bf_i(x)$ contains only messages that will be delivered in transitions of x that come after configuration ρ_i . Possibly $i = glob_M(x, u)$, and in that case the message will be delivered in the transition immediately after configuration ρ_i , which is transition i (see also below).

Transitions So far we have obtained a sequence of configurations ρ_0, ρ_1, ρ_2 , etc. Now we define a sequence of tuples, one tuple per ordinal $i \in \mathbb{N}$, that represents the transition i . Let $i \in \mathbb{N}$. Recall from above that (x_i, s_i) is the unique pair in $\mathcal{N} \times \mathbb{N}$ such that $glob_M(x_i, s_i) = i$. The tuple τ_i is defined as $(\rho_i, x_i, m_i, i, \rho_{i+1})$, where

$$m_i = \{(glob_M(y, t), R(\bar{a})) \mid \text{chosen}_R(y, t, z, u, \bar{a}) \in M, glob_M(z, u) = i\}.$$

Intuitively, m_i selects all messages that arrive in transition i . And since $glob_M(z, u) = i$ implies $z = x_i$ and $u = s_i$, we thus select all messages destined for step s_i of node x_i .

Trace We can show that sequence \mathcal{R} is indeed a legal run of \mathcal{P} on input H such that $trace(\mathcal{R}) = M|_{sch(\mathcal{P})^{LT}}$; this proof can be found in Appendix B of the online appendix to the paper. In the following subsection we show that \mathcal{R} is also fair.

5.3.3 Fair Run

Let \mathcal{R} be the run as constructed in the previous subsection. We now show that \mathcal{R} is fair. For each transition index $i \in \mathbb{N}$, let $\rho_i = (st_i, bf_i)$ denote the source configuration of transition i . Recall from Section 5.1.4 that we have to check two fairness conditions:

1. every node is the active node in an infinite number of transitions; and,
2. for every transition $i \in \mathbb{N}$, for every $y \in \mathcal{N}$, for every pair $(j, \mathbf{f}) \in bf_i(y)$, there is a transition k with $i \leq k$ in which (j, \mathbf{f}) is delivered to y .

We show that \mathcal{R} satisfies the first fairness condition. Let $x \in \mathcal{N}$ be a node, and let $s \in \mathbb{N}$ be a timestamp of x . Consider transition $i = glob_M(x, s)$. This transition has active node $x_i = x$. We can find such a transition with active node x for every timestamp $s \in \mathbb{N}$ of x , and these transitions are all unique because function $glob_M(\cdot)$ is injective. So, there are an infinite number of transitions in \mathcal{R} with active node x .

We show that \mathcal{R} satisfies the second fairness condition. Let $i \in \mathbb{N}$, $y \in \mathcal{N}$, and $(j, \mathbf{f}) \in bf_i(y)$. Denote $\mathbf{f} = R(\bar{a})$. From its construction, the pair $(j, \mathbf{f}) \in bf_i(y)$ implies there are values $x \in \mathcal{N}$, $s \in \mathbb{N}$ and $t \in \mathbb{N}$ such that $\text{chosen}_R(x, s, y, t, \bar{a}) \in M$ and $j = glob_M(x, s) < i \leq glob_M(y, t)$. Denote $k = glob_M(y, t)$. Hence, $i \leq k$ and $(j, \mathbf{f}) \in m_k$ by definition of m_k . Thus (j, \mathbf{f}) is delivered to $x_k = y$ in transition k .

6 Discussion

We have represented distributed programs in Datalog under the stable model semantics. Moreover, we have shown that the stable models represent the desired behavior of the distributed program, as found in a realistic operational semantics. We now discuss some points for future work.

As mentioned, many Datalog-inspired languages have been proposed to implement distributed applications (Loo et al. 2009; Navarro and Rybalchenko 2009; Grumbach and Wang 2010; Abiteboul et al. 2011), and they contain several powerful features such as aggregation and non-determinism (choice). Our current framework already represents the essential features that all these languages possess: reasoning about distributed state and representing message sending. Nonetheless, we have probably not yet explored the full power of stable models. We therefore expect that this work can be extended to languages that incorporate more powerful language constructs such as the ones mentioned above. It might also be possible to remove the syntactic stratification condition that we have used for the deductive rules.

More related to multi-agent systems (Leite et al. 2002; Nigam and Leite 2006; Leite and Soares 2007), it might be interesting to allow logic programs used in declarative networking to dynamically modify their rules. The question would be how (and if) this can be represented in our model-based semantics.

The effect of variants of the model-based semantics can be studied. For example, messages can be sent into the past when the causality rules are removed. Then, one might ask which (classes of) programs still work “correctly” under such a non-causal semantics; some preliminary results are in (Ameloot and den Bussche 2014).

Lastly, we can think about the output of distributed Datalog programs. Marczak et al. (2011) define the output with *ultimate* facts, which are facts that will eventually always be present on the network. This way, the output of a run (or equivalently stable model) can be defined. Then, a *consistent* program is required to produce the same output in every run. For consistent programs, the output on an input distributed database instance can thus be defined as the output of any run. We can now consider the following decision problem: for a consistent program, an input distributed database instance for that program, and a fact, decide if this fact is output by the program on that input. We think that decidability depends on the semantics of the message buffers. In this paper, we have represented per addressee duplicate messages in its message buffer. This is a realistic representation, since in a real network, the same message can be sent multiple times, and hence, multiple instances of the same message can be in transmission simultaneously. If we would forbid duplicate messages in the buffers, then the decision problem becomes decidable because only a finite number of configurations would be possible by finiteness of the input domain. But when duplicates are preserved, the number of configurations is not limited, and we expect that the problem will be undecidable in general. However, we might want to investigate whether decidability can be obtained in particular (syntactically defined) cases. If so, it might be interesting for those cases to find finite representations of the stable models. This could serve as a more intuitive

programmer abstraction, or it could perhaps be used to more efficiently simulate the behavior of the network for testing purposes.

Acknowledgment

The second author thanks Serge Abiteboul for a number of interesting discussions.

References

- ABITEBOUL, S., BIENVENU, M., GALLAND, A., ET AL. 2011. A rule-based language for Web data management. In *Proceedings 30th ACM Symposium on Principles of Database Systems*, pp. 293–304. ACM Press.
- ABITEBOUL, S., HULL, R., AND VIANU, V. 1995. *Foundations of Databases*. Addison-Wesley.
- ABRIAL, J. 2010. *Modeling in Event-B – System and Software Engineering*. Cambridge University Press.
- ALFERES, J., PEREIRA, L., PRZYMUSINSKA, H., AND PRZYMUSINSKI, T. 2002. LUPS—a language for updating logic programs. *Artificial Intelligence* 138, 1–2, 87–116.
- ALVARO, P., CONDIE, T., CONWAY, N., HELLERSTEIN, J., AND SEARS, R. 2009. I do declare: consensus in a logic language. *Operating Systems Review* 43, 4, 25–30.
- ALVARO, P., CONWAY, N., HELLERSTEIN, J., AND MARCZAK, W. 2011. Consistency analysis in Bloom: A CALM and collected approach. In *Proceedings 5th Biennial Conference on Innovative Data Systems Research*, pp. 249–260. www.cidrdb.org.
- ALVARO, P., MARCZAK, W., ET AL. 2009. Dedalus: Datalog in time and space. Technical Report EECS-2009-173, University of California, Berkeley.
- ALVARO, P., MARCZAK, W., ET AL. 2011. Dedalus: Datalog in time and space. See de Moor et al. (2011), pp. 262–281.
- AMELOOT, T. AND DEN BUSSCHE, J. V. 2014. Positive Dedalus programs tolerate non-causality. *Journal of Computer and System Sciences* 80, 7, 1191–1213.
- AMELOOT, T., NEVEN, F., AND VAN DEN BUSSCHE, J. 2011. Relational transducers for declarative networking. In *Proceedings 30th ACM Symposium on Principles of Database Systems*, pp. 283–292. ACM Press.
- AMELOOT, T., NEVEN, F., AND VAN DEN BUSSCHE, J. 2013. Relational transducers for declarative networking. *Journal of the ACM* 60, 2, 15:1–15:38.
- AMELOOT, T. AND VAN DEN BUSSCHE, J. 2012. Deciding eventual consistency for a simple class of relational transducer networks. In *Proceedings of the 15th International Conference on Database Theory*, pp. 86–98. ACM Press.
- APT, K. AND BOL, R. 1994. Logic programming and negation: A survey. *The Journal of Logic Programming* 19-20, Supplement 1, 0, 9–71.
- APT, K., FRANCEZ, N., AND KATZ, S. 1988. Appraising fairness in languages for distributed programming. *Distributed Computing* 2, 226–241.
- ATTIYA, H. AND WELCH, J. 2004. *Distributed Computing: Fundamentals, Simulations, and Advanced Topics*. Wiley.
- CAVAGE, M. 2013. There’s just no getting around it: You’re building a distributed system. *ACM Queue* 11, 4.
- DE MOOR, O., GOTTLÖB, G., FURCHE, T., AND SELLERS, A. (Eds.) 2011. *Datalog Reloaded: First International Workshop, Datalog 2010*, Volume 6702 of *Lecture Notes in Computer Science*.

- DEUTSCH, A., SUI, L., VIANU, V., AND ZHOU, D. 2006. Verification of communicating data-driven Web services. In *Proceedings 25th ACM Symposium on Principles of Database Systems*, pp. 90–99. ACM Press.
- FRANCEZ, N. 1986. *Fairness*. Springer-Verlag New York, Inc., New York, NY, USA.
- GELFOND, M. AND LIFSCHITZ, V. 1988. The stable model semantics for logic programming. In *Proceedings of the Fifth International Conference on Logic Programming*, pp. 1070–1080. MIT Press.
- GRUMBACH, S. AND WANG, F. 2010. Netlog, a rule-based language for distributed programming. In M. CARRO AND R. PEÑA (Eds.), *Proceedings 12th International Symposium on Practical Aspects of Declarative Languages*, Volume 5937 of *Lecture Notes in Computer Science*, pp. 88–103.
- HELLERSTEIN, J. 2010a. Datalog redux: experience and conjecture. Video available (under the title “The Declarative Imperative”) from <http://db.cs.berkeley.edu/jmh/>. PODS 2010 keynote.
- HELLERSTEIN, J. 2010b. The declarative imperative: experiences and conjectures in distributed logic. *SIGMOD Record* 39, 1, 5–19.
- HUANG, S., GREEN, T., AND LOO, B. 2011. Datalog and emerging applications: an interactive tutorial. In *Proceedings of the 2011 ACM SIGMOD International Conference on the Management of Data*, SIGMOD ’11, pp. 1213–1216. ACM.
- INTERLANDI, M., TANCA, L., AND BERGAMASCHI, S. 2013. Datalog in time and space, synchronously. In *7th Alberto Mendelzon International Workshop on Foundations of Data Management*.
- JIM, T. 2001. SD3: A trust management system with certified evaluation. In *Proceedings of the 2001 IEEE Symposium on Security and Privacy*, SP, pp. 106–115. IEEE Computer Society.
- KRISHNAMURTHY, R. AND NAQVI, S. 1988. Non-deterministic choice in Datalog. In *Proceedings of the Third International Conference on Data and Knowledge Bases*, pp. 416–424.
- LAMPORT, L. 2000a. Distributed algorithms in TLA (abstract). In *Proceedings of the nineteenth ACM symposium on Principles of distributed computing*, pp. 3. ACM Press.
- LAMPORT, L. 2000b. Fairness and hyperfairness. *Distributed Computing* 13, 239–245.
- LEITE, J., ALFERES, J., AND PEREIRA, L. 2002. Minerva – a dynamic logic programming agent architecture. In *Revised Papers from the 8th International Workshop on Intelligent Agents VIII*, ATAL, pp. 141–157. Springer-Verlag.
- LEITE, J. AND SOARES, L. 2007. Adding evolving abilities to a multi-agent system. In *Proceedings of the 7th International Conference on Computational Logic in Multi-agent Systems*, CLIMA VII’06, pp. 246–265. Springer-Verlag.
- LOBO, J., MA, J., RUSSO, A., AND LE, F. 2012. Declarative distributed computing. In E. ERDEM, J. LEE, Y. LIERLER, AND D. PEARCE (Eds.), *Correct Reasoning - Essays on Logic-Based AI in Honour of Vladimir Lifschitz*, Volume 7265 of *Lecture Notes in Computer Science*, pp. 454–470. Springer.
- LOO, B. ET AL. 2009. Declarative networking. *Commun. ACM* 52, 11, 87–95.
- LYNCH, N. 1996. *Distributed Algorithms*. Morgan Kaufmann.
- MA, J., LE, F., WOOD, D., RUSSO, A., AND LOBO, J. 2013. A declarative approach to distributed computing: Specification, execution and analysis. *Theory and Practice of Logic Programming* 13, 815–830.
- MARCZAK, W., ALVARO, P., CONWAY, N., HELLERSTEIN, J., AND MAIER, D. 2011. Confluence analysis for distributed programs: A model-theoretic approach. Technical Report UCB/EECS-2011-154 (Dec), EECS Department, University of California, Berkeley.

- MAREK, V. AND TRUSZCZYNSKI, M. 1999. Stable models and an alternative logic programming paradigm. In K. APT, V. MAREK, M. TRUSZCZYNSKI, AND D. WARREN (Eds.), *The Logic Programming Paradigm*, Artificial Intelligence, pp. 375–398. Springer Berlin Heidelberg.
- NAVARRO, J. AND RYBALCHENKO, A. 2009. Operational semantics for declarative networking. In A. GILL AND T. SWIFT (Eds.), *Proceedings 11th International Symposium on Practical Aspects of Declarative Languages*, Volume 5419 of *Lecture Notes in Computer Science*, pp. 76–90.
- NIGAM, V. AND LEITE, J. 2006. A dynamic logic programming based system for agents with declarative goals. In *Proceedings of the 4th International Conference on Declarative Agent Languages and Technologies*, DALT, pp. 174–190. Springer-Verlag.
- SACCÀ, D. AND ZANIOLO, C. 1990. Stable models and non-determinism in logic programs with negation. In *Proceedings of the Ninth ACM Symposium on Principles of Database Systems*, pp. 205–217. ACM Press.
- VARDI, M. 1982. The complexity of relational query languages. In *Proceedings 14th ACM Symposium on the Theory of Computing*, pp. 137–146.
- WOODCOCK, J. AND DAVIES, J. 1996. *Using Z: Specification, Refinement, and Proof*. Prentice Hall.
- ZHANG, Q., CHENG, L., AND BOUTABA, R. 2010. Cloud computing: state-of-the-art and research challenges. *Journal of Internet Services and Applications* 1, 7–18.
- ZINN, D., GREEN, T., AND LUDAESCHER, B. 2012. Win-move is coordination-free (sometimes). In *Proceedings of the 15th International Conference on Database Theory*, pp. 99–113. ACM Press.

Appendix

General Remarks

Let \mathcal{P} be a Dedalus program. Recall from Section 5.1.2 that $\text{deduc}_{\mathcal{P}} \subseteq \mathcal{P}$ is the subset of all (unmodified) deductive rules. The semantics of $\text{deduc}_{\mathcal{P}}$ is given by the stratified semantics. Although the semantics of $\text{deduc}_{\mathcal{P}}$ does not depend on the chosen syntactic stratification, for technical convenience in the proofs, we will fix an arbitrary syntactic stratification for $\text{deduc}_{\mathcal{P}}$. Whenever we refer to the stratum number of an *idb* relation, we implicitly use this fixed syntactic stratification. Stratum numbers start at 1.

Appendix A Run to Model: Proof Details

In the context of Section 5.2.2, we show that M is a model of \mathcal{P} on input H . Let G abbreviate the ground program $\text{ground}_M(C, I)$, where $C = \text{pure}(\mathcal{P})$ and $I = \text{decl}(H)$. To show that M is a stable model, we have to show $M = N$ where $N = G(\text{decl}(H))$. The inclusions $M \subseteq N$ and $N \subseteq M$ are shown respectively in Sections A.1 and A.2. We use the notations of Section 5.2.2.

A.1 Inclusion $M \subseteq N$

By definition,

$$M = \text{decl}(H) \cup \bigcup_{i \in \mathbb{N}} \text{trans}_{\mathcal{R}}^{[i]}.$$

We immediately have $\text{decl}(H) \subseteq N$ by the semantics of G . Next, we define for uniformity the set $\text{trans}_{\mathcal{R}}^{[-1]} = \emptyset$. We will show by induction on $i = -1, 0, 1, \dots$, that $\text{trans}_{\mathcal{R}}^{[i]} \subseteq N$. The base case ($i = -1$) is clear. For the induction hypothesis, let $i \geq 0$, and assume for all $j \in \{-1, 0, \dots, i-1\}$ that $\text{trans}_{\mathcal{R}}^{[j]} \subseteq N$. We show that $\text{trans}_{\mathcal{R}}^{[i]} \subseteq N$. By definition,

$$\text{trans}_{\mathcal{R}}^{[i]} = \text{caus}_{\mathcal{R}}^{[i]} \cup \text{fin}_{\mathcal{R}}^{[i]} \cup \text{duc}_{\mathcal{R}}^{[i]} \cup \text{snd}_{\mathcal{R}}^{[i]}.$$

We show inclusion of these four sets in N below. Auxiliary claims can be found in Section A.1.5.

A.1.1 Causality

We show that $\text{caus}_{\mathcal{R}}^{[i]} \subseteq N$. Concretely, let $(x, s) \in \mathcal{N} \times \mathbb{N}$ such that $(x, s) \prec_{\mathcal{R}} (x_i, s_i)$. We show $\text{before}(x, s, x_i, s_i) \in N$. We distinguish between the following cases.

Local edge Suppose $(x, s) \prec_{\mathcal{R}} (x_i, s_i)$ is a local edge, i.e., $x = x_i$ and $s_i = s + 1$. Because rule (7) is positive, the following ground rule is always in G :

$$\text{before}(x, s, x, s + 1) \leftarrow \text{all}(x), \text{tsucc}(s, s + 1).$$

The body facts of this ground rule are in $\text{decl}(H) \subseteq N$; hence, the rule derives $\text{before}(x, s, x, s + 1) = \text{before}(x, s, x_i, s_i) \in N$.

Message edge Suppose $(x, s) \prec_{\mathcal{R}} (x_i, s_i)$ is a message edge, i.e., there is an earlier transition $j < i$ with $j = \text{glob}_{\mathcal{R}}(x, s)$, in which x sends a message \mathbf{f} to x_i such that $\alpha_{\mathcal{R}}(j, x_i, \mathbf{f}) = i$. Denote $\mathbf{f} = R(\bar{a})$. Because rules of the form (10) in $\text{pure}(\mathcal{P})$ are positive, the following ground rule is always in G :

$$\text{before}(x, s, x_i, s_i) \leftarrow \text{chosen}_R(x, s, x_i, s_i, \bar{a}).$$

We show $\text{chosen}_R(x, s, x_i, s_i, \bar{a}) \in N$, so that $\text{before}(x, s, x_i, s_i) \in N$, as desired. Since $j = \text{glob}_{\mathcal{R}}(x, s)$, we have $x_j = x$ and $s_j = s$. Also using $s_i = \text{loc}_{\mathcal{R}}(i)$, we have

$$\text{chosen}_R(x, s, x_i, s_i, \bar{a}) \in \text{snd}_{\mathcal{R}}^{[j]} \subseteq \text{trans}_{\mathcal{R}}^{[j]}.$$

Lastly, we have $\text{trans}_{\mathcal{R}}^{[j]} \subseteq N$ by applying the induction hypothesis.

Transitive edge Suppose $(x, s) \prec_{\mathcal{R}} (x_i, s_i)$ is not a local edge nor a message edge. Then we can choose a pair $(z, u) \in \mathcal{N} \times \mathbb{N}$ such that $(x, s) \prec_{\mathcal{R}} (z, u)$ and $(z, u) \prec_{\mathcal{R}} (x_i, s_i)$, but also such that $(z, u) \prec_{\mathcal{R}} (x_i, s_i)$ is a local edge or a message edge. Because rule (8) is positive, the following ground rule is always in G :

$$\text{before}(x, s, x_i, s_i) \leftarrow \text{before}(x, s, z, u), \text{before}(z, u, x_i, s_i).$$

We now show that the body of this rule is in N , so that $\text{before}(x, s, x_i, s_i) \in N$, as desired. Denote $j = \text{glob}_{\mathcal{R}}(z, u)$. First, because $(x, s) \prec_{\mathcal{R}} (z, u)$, we have $\text{before}(x, s, z, u) \in \text{caus}_{\mathcal{R}}^{[j]}$. Next, because $(z, u) \prec_{\mathcal{R}} (x_i, s_i)$, we have $j < i$ by Lemma 1. So, by applying the induction hypothesis to j , we have $\text{before}(x, s, z, u) \in N$. Secondly, because $(z, u) \prec_{\mathcal{R}} (x_i, s_i)$ is a local edge or a message edge, we have $\text{before}(z, u, x_i, s_i) \in N$ as shown in the preceding two cases.

A.1.2 Finite Messages

We show that $\text{fin}_{\mathcal{R}}^{[i]} \subseteq N$. Let $\text{senders}_{\mathcal{R}}^{[i]}$ be as defined in Section 5.2.2. For each of the different kinds of facts in $\text{fin}_{\mathcal{R}}^{[i]}$, we show inclusion in N .

Senders Let $\text{hasSender}(x_i, s_i, x, s) \in \text{fin}_{\mathcal{R}}^{[i]}$. We have $(x, s) \in \text{senders}_{\mathcal{R}}^{[i]}$, which means that x during step s sends some message fact $R(\bar{a})$ that arrives in step s_i of x_i . Rules in $\text{pure}(\mathcal{P})$ of the form (11) have a negative rcvInf -atom in their body. But since we have not added any rcvInf -facts to M , including $\text{rcvInf}(x_i, s_i)$, the following rule is in G :

$$\text{hasSender}(x_i, s_i, x, s) \leftarrow \text{chosen}_R(x, s, x_i, s_i, \bar{a}).$$

We are left to show that $\text{chosen}_R(x, s, x_i, s_i, \bar{a}) \in N$. Denote $j = \text{glob}_{\mathcal{R}}(x, s)$. Using that $x = x_j$ and $s = s_j$, we have $\text{chosen}_R(x, s, x_i, s_i, \bar{a}) \in \text{snd}_{\mathcal{R}}^{[j]}$. Because $j < i$ by the operational semantics, we can apply the induction hypothesis to j to know $\text{snd}_{\mathcal{R}}^{[j]} \subseteq N$.

Comparison of timestamps Let $\text{isSmaller}(x_i, s_i, x, s) \in \text{fin}_{\mathcal{R}}^{[i]}$. We have $(x, s) \in \text{senders}_{\mathcal{R}}^{[i]}$ and there is a timestamp $s' \in \mathbb{N}$ so that $(x, s') \in \text{senders}_{\mathcal{R}}^{[i]}$ and $s < s'$. Rule (12) is positive and therefore the following ground rule is always in G :

$$\begin{aligned} \text{isSmaller}(x_i, s_i, x, s) \leftarrow & \text{hasSender}(x_i, s_i, x, s), \text{hasSender}(x_i, s_i, x, s'), \\ & s < s'. \end{aligned}$$

We immediately have $(s < s') \in \text{decl}(H) \subseteq N$. By construction of $\text{fin}_{\mathcal{R}}^{[i]}$, we also have $\text{hasSender}(x_i, s_i, x, s) \in \text{fin}_{\mathcal{R}}^{[i]}$ and $\text{hasSender}(x_i, s_i, x, s') \in \text{fin}_{\mathcal{R}}^{[i]}$, and thus both facts are also in N as shown above. Hence the previous ground rule derives $\text{isSmaller}(x_i, s_i, x, s) \in N$.

Maximum timestamp Let $\text{hasMax}(x_i, s_i, x) \in \text{fin}_{\mathcal{R}}^{[i]}$. Thus x is a sender-node mentioned in $\text{senders}_{\mathcal{R}}^{[i]}$. Let s be the maximum send-timestamp of x in $\text{senders}_{\mathcal{R}}^{[i]}$, which surely exists because $\text{senders}_{\mathcal{R}}^{[i]}$ is finite. We have not added $\text{isSmaller}(x_i, s_i, x, s)$ to $\text{fin}_{\mathcal{R}}^{[i]}$, and thus also not to M . Although rule (13) contains a negated isSmaller -atom, $\text{isSmaller}(x_i, s_i, x, s) \notin M$ implies that the following ground rule is in G :

$$\text{hasMax}(x_i, s_i, x) \leftarrow \text{hasSender}(x_i, s_i, x, s).$$

Moreover, $(x, s) \in \text{senders}_{\mathcal{R}}^{[i]}$ implies $\text{hasSender}(x_i, s_i, x, s) \in N$, and thus the previous ground rule derives $\text{hasMax}(x_i, s_i, x) \in N$, as desired.

A.1.3 Deductive

We show that $\text{duc}_{\mathcal{R}}^{[i]} \subseteq N$. By definition, $\text{duc}_{\mathcal{R}}^{[i]} = D_i^{\uparrow x_i, s_i}$, where D_i is the output of subprogram $\text{deduc}_{\mathcal{P}}$ during transition i . Recall from Section 5.1.3 that $\text{deduc}_{\mathcal{P}}$ is given the following input during transition i :

$$\text{st}_i(x_i) \cup \text{untag}(m_i),$$

where st_i denotes the state at the beginning of transition i , and m_i is the set of (tagged) messages delivered during transition i . If we can show that $(st_i(x_i) \cup \text{untag}(m_i))^{\uparrow x_i, s_i} \subseteq N$, then we can apply Claim 8 to know that $D_i^{\uparrow x_i, s_i} \subseteq N$, as desired.

State We first show $st_i(x_i)^{\uparrow x_i, s_i} \subseteq N$. There are two cases:

- Suppose $s_i = 0$, i.e., i is the first transition of \mathcal{R} with active node x_i . Then $st_i(x_i) = H(x_i)$ by the operational semantics, which gives $st_i(x_i)^{\uparrow x_i, s_i} \subseteq \text{decl}(H) \subseteq N$ by definition of $\text{decl}(H)$.
- Suppose $s_i > 0$. Then we can consider the last transition j of x_i that came before i . By the operational semantics, we have $st_i(x_i) = st_{j+1}(x_i)$, where st_{j+1} is the state resulting from transition j . More concretely, $st_i(x_i) = H(x_i) \cup \text{induc}_{\mathcal{P}}(D_j)$, with D_j the output of $\text{deduc}_{\mathcal{P}}$ during transition j . As in the previous case, we already know $H(x_i)^{\uparrow x_i, s_i} \subseteq \text{decl}(H)$. Now, by applying the induction hypothesis to j , we have $\text{duc}_{\mathcal{R}}^{[j]} \subseteq \text{trans}_{\mathcal{R}}^{[j]} \subseteq N$. Next, by applying Claim 10, and by using $s_i = s_j + 1$, we obtain

$$\begin{aligned} st_i(x_i)^{\uparrow x_i, s_i} &= H(x_i)^{\uparrow x_i, s_i} \cup \text{induc}_{\mathcal{P}}(D_j)^{\uparrow x_i, s_j+1} \\ &\subseteq N. \end{aligned}$$

Messages Now we show $\text{untag}(m_i)^{\uparrow x_i, s_i} \subseteq N$. Let $\mathbf{f} \in \text{untag}(m_i)$. We have to show that $\mathbf{f}^{\uparrow x_i, s_i} \in N$. First, because $\mathbf{f} \in \text{untag}(m_i)$, there is a transition k with $k < i$ such that $(k, \mathbf{f}) \in m_i$, i.e., the fact \mathbf{f} was sent to x_i during transition k (by node x_k). Denote $\mathbf{f} = R(\bar{a})$. So, there must be an asynchronous rule with head-predicate R in \mathcal{P} , which has a corresponding rule in $\text{pure}(\mathcal{P})$ of the form (6). Rules of the form (6) are positive and thus the following ground rule is always in G :

$$R(x_i, s_i, \bar{a}) \leftarrow \text{chosen}_R(x_k, s_k, x_i, s_i, \bar{a}).$$

We show $\text{chosen}_R(x_k, s_k, x_i, s_i, \bar{a}) \in N$, so that the rule derives $\mathbf{f}^{\uparrow x_i, s_i} \in N$, as desired. Because x_k sends \mathbf{f} to x_i during transition k , and i is the transition in which this message is delivered to x_i , we have $\text{chosen}_R(x_k, s_k, x_i, s_i, \bar{a}) \in \text{snd}_{\mathcal{R}}^{[k]} \subseteq \text{trans}_{\mathcal{R}}^{[k]}$. By applying the induction hypothesis to k , we have $\text{snd}_{\mathcal{R}}^{[k]} \subseteq N$.

A.1.4 Sending

We show that $\text{snd}_{\mathcal{R}}^{[i]} \subseteq N$. For each kind of fact in $\text{snd}_{\mathcal{R}}^{[i]}$ we show inclusion in N .

Candidates Let $\text{cand}_R(x_i, s_i, y, t, \bar{a}) \in \text{snd}_{\mathcal{R}}^{[i]}$. We have $R(y, \bar{a}) \in \text{mesg}_{\mathcal{R}}^{[i]}$, $t \in \mathbb{N}$ and $(y, t) \not\prec_{\mathcal{R}} (x_i, s_i)$. Since $D_i^{\uparrow x_i, s_i} \subseteq N$ (see above), we can use Claim 11 to obtain $\text{cand}_R(x_i, s_i, y, t, \bar{a}) \in N$, as desired.

Chosen Let $\text{chosen}_R(x_i, s_i, y, t, \bar{a}) \in \text{snd}_{\mathcal{R}}^{[i]}$. We have $R(y, \bar{a}) \in \text{mesg}_{\mathcal{R}}^{[i]}$ and $t = \text{loc}_{\mathcal{R}}(j)$ with $j = \alpha_{\mathcal{R}}(i, y, R(\bar{a}))$. Because $R(y, \bar{a}) \in \text{mesg}_{\mathcal{R}}^{[i]}$, this fact was produced by $\text{async}_{\mathcal{P}}$, and thus there is an asynchronous rule in \mathcal{P} with head-predicate R .

This asynchronous rule has a corresponding rule in $\text{pure}(\mathcal{P})$ of the form (4), that contains a negated **other**_R-atom in the body. But by construction of $\text{snd}_{\mathcal{R}}^{[i]}$, we have not added **other**_R(x_i, s_i, y, t, \bar{a}) to $\text{snd}_{\mathcal{R}}^{[i]}$, and thus also not to M . Therefore the following ground rule of the form (4) is in G :

$$\text{chosen}_R(x_i, s_i, y, t, \bar{a}) \leftarrow \text{cand}_R(x_i, s_i, y, t, \bar{a}).$$

Because $j > i$ by the operational semantics, we have $(y, t) \not\prec_{\mathcal{R}} (x_i, s_i)$ by Lemma 1. Thus, by construction of $\text{snd}_{\mathcal{R}}^{[i]}$, we have $\text{cand}_R(x_i, s_i, y, t, \bar{a}) \in \text{snd}_{\mathcal{R}}^{[i]}$, in which case $\text{cand}_R(x_i, s_i, y, t, \bar{a}) \in N$ (shown above). Hence, the previous ground rule derives $\text{chosen}_R(x_i, s_i, y, t, \bar{a}) \in N$, as desired.

Other Let $R(y, \bar{a})$ and t be from above. Let **other**_R(x_i, s_i, y, u, \bar{a}) $\in \text{snd}_{\mathcal{R}}^{[i]}$. We have $u \in \mathbb{N}$, $(y, u) \not\prec_{\mathcal{R}} (x_i, s_i)$ and $u \neq t$. Because rule (5) is positive, the following ground rule is in G :

$$\begin{aligned} \text{other}_R(x_i, s_i, y, u, \bar{a}) \leftarrow & \text{cand}_R(x_i, s_i, y, u, \bar{a}), \text{chosen}_R(x_i, s_i, y, t, \bar{a}), \\ & u \neq t. \end{aligned}$$

We immediately have $(u \neq t) \in \text{decl}(H) \subseteq N$. Now we show that the other body facts are in N , so the rule derives **other**_R(x_i, s_i, y, u, \bar{a}) $\in N$, as desired. Because $(y, u) \not\prec_{\mathcal{R}} (x_i, s_i)$, by construction of $\text{snd}_{\mathcal{R}}^{[i]}$, we have $\text{cand}_R(x_i, s_i, y, u, \bar{a}) \in \text{snd}_{\mathcal{R}}^{[i]}$ and thus $\text{cand}_R(x_i, s_i, y, u, \bar{a}) \in N$ (shown above). Moreover, it was shown above that $\text{chosen}_R(x_i, s_i, y, t, \bar{a}) \in N$.

A.1.5 Subclaims

Claim 8

Let i be a transition of \mathcal{R} . If $(st_i(x_i) \cup \text{untag}(m_i))^{\uparrow x_i, s_i} \subseteq N$, then $D_i^{\uparrow x_i, s_i} \subseteq N$.

Proof

Abbreviate $I_i = st_i(x_i) \cup \text{untag}(m_i)$. Recall that $D_i = \text{deduc}_{\mathcal{P}}(I_i)$, which is computed with the stratified semantics.

For $k \in \mathbb{N}$, we write $D_i^{\rightarrow k}$ to denote the set obtained by adding to I_i all facts derived in stratum 1 up to stratum k during the computation of D_i . For the largest stratum number n of $\text{deduc}_{\mathcal{P}}$, we have $D_i^{\rightarrow n} = D_i$. Also, because stratum numbers start at 1, we have $D_i^{\rightarrow 0} = I_i$. We show by induction on $k = 0, 1, 2, \dots, n$, that $(D_i^{\rightarrow k})^{\uparrow x_i, s_i} \subseteq N$.

Base case For the base case, $k = 0$, the property holds by the given assumption $I_i^{\uparrow x_i, s_i} \subseteq N$.

Induction hypothesis For the induction hypothesis, assume for some stratum number k with $k \geq 1$ that $(D_i^{\rightarrow k-1})^{\uparrow x_i, s_i} \subseteq N$.

Inductive step For the inductive step, we show that $(D_i^{\rightarrow k})^{\uparrow x_i, s_i} \subseteq N$. Recall that the input of stratum k in $deduc_{\mathcal{P}}$ is the set $D_i^{\rightarrow k-1}$, and the semantics is given by the fixpoint semantics of semi-positive Datalog⁻ (see Section 3.2.2). So, we can consider $D_i^{\rightarrow k}$ to be a fixpoint, i.e., as the set $\bigcup_{l \in \mathbb{N}} A_l$ with $A_0 = D_i^{\rightarrow k-1}$ and $A_l = T(A_{l-1})$ for each $l \geq 1$, where T is the immediate consequence operator of stratum k . We show by inner induction on $l = 0, 1$, etc, that

$$(A_l)^{\uparrow x_i, s_i} \subseteq N.$$

For the base case ($l = 0$), we have $A_0 = D_i^{\rightarrow k-1}$, for which we can apply the outer induction hypothesis to know that $(D_i^{\rightarrow k-1})^{\uparrow x_i, s_i} = (A_0)^{\uparrow x_i, s_i} \subseteq N$, as desired. For the inner induction hypothesis, we assume for some $l \geq 1$ that $(A_{l-1})^{\uparrow x_i, s_i} \subseteq N$. For the inner inductive step, we show that $(A_l)^{\uparrow x_i, s_i} \subseteq N$. Let $\mathbf{f} \in A_l \setminus A_{l-1}$. Let $\varphi \in deduc_{\mathcal{P}}$ and V be a rule from stratum k and valuation respectively that have derived \mathbf{f} . Let φ' be the rule in $pure(\mathcal{P})$ obtained by applying the transformation (1) to φ . Let V' be V extended to assign x_i and s_i to the new variables in φ' that represent the location and timestamp respectively. Note in particular that $V'(pos_{\varphi'}) = V(pos_{\varphi})^{\uparrow x_i, s_i}$ and $V'(neg_{\varphi'}) = V(neg_{\varphi})^{\uparrow x_i, s_i}$. Let ψ be the *positive* ground rule obtained by applying V' to φ' and by subsequently removing all negative (ground) body atoms. We show that $\psi \in G$ and that its body is in N , so that ψ derives $head_{\psi} = \mathbf{f}^{\uparrow x_i, s_i} \in N$, as desired.

- In order for ψ to be in G , it is required that $V'(neg_{\varphi'}) \cap M = \emptyset$. Because V is satisfying for φ , and negation in φ is only applied to lower strata, we have $V(neg_{\varphi}) \cap D_i^{\rightarrow k-1} = \emptyset$. Moreover, since a relation is computed in only one stratum of $deduc_{\mathcal{P}}$, we overall have $V(neg_{\varphi}) \cap D_i = \emptyset$. Then by Claim 9 we have $V(neg_{\varphi})^{\uparrow x_i, s_i} \cap M = \emptyset$. Hence,

$$V'(neg_{\varphi'}) \cap M = \emptyset.$$

- Now we show that $pos_{\psi} \subseteq N$. Because V is satisfying for φ , we have $V(pos_{\varphi}) \subseteq A_{l-1}$, and by applying the inner induction hypothesis we have $V(pos_{\varphi})^{\uparrow x_i, s_i} \subseteq N$. Therefore, $pos_{\psi} = V'(pos_{\varphi'}) \subseteq N$.

□

Claim 9

Let i be a transition of \mathcal{R} . Let I be a set of facts over $sch(\mathcal{P})$. If $I \cap D_i = \emptyset$ then $I^{\uparrow x_i, s_i} \cap M = \emptyset$.

Proof

If a fact $\mathbf{f} \in M$ is over schema $sch(\mathcal{P})^{\text{LT}}$ and has location specifier x_i and timestamp s_i then $\mathbf{f} \in \text{duc}_{\mathcal{R}}^{[i]}$ because (i) for any transition j there are no facts over $sch(\mathcal{P})^{\text{LT}}$ in $\text{caus}_{\mathcal{R}}^{[j]}$, $\text{fin}_{\mathcal{R}}^{[j]}$ or $\text{snd}_{\mathcal{R}}^{[j]}$; (ii) we only add facts with location specifier x_i to $\text{duc}_{\mathcal{R}}^{[j]}$ if j is a transition of node x_i ; and, (iii) for every transition j of node x_i , if $i \neq j$ then $\text{loc}_{\mathcal{R}}(j) \neq s_i$.

Hence, it suffices to show $I^{\uparrow x_i, s_i} \cap \text{duc}_{\mathcal{R}}^{[i]} = \emptyset$. But this is immediate from $I \cap D_i = \emptyset$ because $\text{duc}_{\mathcal{R}}^{[i]}$ equals $D_i^{\uparrow x_i, s_i}$ by definition. □

Claim 10

Let j be a transition of \mathcal{R} . Let D_j be the output of $\text{deduc}_{\mathcal{P}}$ during transition j . Suppose $\text{duc}_{\mathcal{R}}^{[j]} \subseteq N$. We have $\text{induc}_{\mathcal{P}}(D_j)^{\uparrow x_j, s_j+1} \subseteq N$.

Proof

Let $\mathbf{f} \in \text{induc}_{\mathcal{P}}(D_j)$. Let $\varphi \in \text{induc}_{\mathcal{P}}$ and V respectively be a rule and valuation that have derived \mathbf{f} . Let φ' be the rule in $\text{pure}(\mathcal{P})$ that is obtained after applying transformation (2) to φ . Thus, besides the additional location variable, the rule φ' has two timestamp variables, one in the body and one in the head. Moreover, the body contains an additional positive tsucc -atom. Let V' be V extended to assign x_j to the location variable, and to assign timestamps s_j and $s_j + 1$ to the body and head timestamp variables respectively. Let ψ be the *positive* ground rule obtained from φ' by applying valuation V' and by subsequently removing all negative (ground) body atoms. We show that $\psi \in G$ and that its body is in N , so that ψ derives $\text{head}_{\psi} = \mathbf{f}^{\uparrow x_j, s_j+1} \in N$, as desired.

- For ψ to be in G , we require $V'(\text{neg}_{\varphi'}) \cap M = \emptyset$. Since $V'(\text{neg}_{\varphi'}) = V(\text{neg}_{\varphi})^{\uparrow x_j, s_j}$, it suffices to show $V(\text{neg}_{\varphi})^{\uparrow x_j, s_j} \cap M = \emptyset$. Because V is satisfying for φ , we have $V(\text{neg}_{\varphi}) \cap D_j = \emptyset$. Then, by Claim 9 we have $V(\text{neg}_{\varphi})^{\uparrow x_j, s_j} \cap M = \emptyset$.
- Now we show $V'(\text{pos}_{\varphi'}) \subseteq N$. The set $V'(\text{pos}_{\varphi'})$ consists of the facts $V(\text{pos}_{\varphi})^{\uparrow x_j, s_j}$ and the fact $\text{tsucc}(s_j, s_j + 1)$. The latter fact is in $\text{decl}(H)$ and thus in N . For the other facts, because V is satisfying for φ , we have $V(\text{pos}_{\varphi}) \subseteq D_j$ and thus $V(\text{pos}_{\varphi})^{\uparrow x_j, s_j} \subseteq D_j^{\uparrow x_j, s_j} = \text{duc}_{\mathcal{R}}^{[j]}$. And by using the given assumption $\text{duc}_{\mathcal{R}}^{[j]} \subseteq N$, we obtain the inclusion in N .

□

Claim 11

Let i be a transition of \mathcal{R} . Suppose $D_i^{\uparrow x_i, s_i} \subseteq N$. For each $R(y, \bar{a}) \in \text{mesg}_{\mathcal{R}}^{[i]}$ and timestamp $t \in \mathbb{N}$ with $(y, t) \not\prec_{\mathcal{R}} (x_i, s_i)$ we have

$$\text{cand}_R(x_i, s_i, y, t, \bar{a}) \in N.$$

Proof

By definition of $\text{mesg}_{\mathcal{R}}^{[i]}$, we have $R(y, \bar{a}) \in \text{async}_{\mathcal{P}}(D_i)$. Let $\varphi \in \text{async}_{\mathcal{P}}$ and V be a rule and valuation that have produced $R(y, \bar{a})$. Let $\varphi' \in \mathcal{P}$ be the original asynchronous rule on which φ is based. Let $\varphi'' \in \text{pure}(\mathcal{P})$ be the rule obtained from φ' by applying transformation (9). Let V'' be valuation V extended to assign x_i and s_i to respectively the sender location and sender timestamp of φ'' , and to assign y and t respectively to the addressee location and addressee arrival timestamp. Let ψ denote the *positive* ground rule that is obtained from φ'' by applying valuation V'' and by subsequently removing all negative (ground) body atoms. We show that $\psi \in G$ and that its body is in N , so that ψ derives $\text{head}_{\psi} = \text{cand}_R(x_i, s_i, y, t, \bar{a}) \in N$, as desired.

- For ψ to be in G , we require $V''(neg_{\varphi''}) \cap M = \emptyset$. By construction of φ'' , the set $V''(neg_{\varphi''})$ consists of the facts $V(neg_{\varphi})^{\uparrow x_i, s_i}$ and the fact **before**(y, t, x_i, s_i). First, because V is satisfying for φ , we have $V(neg_{\varphi}) \cap D_i = \emptyset$, and thus $V(neg_{\varphi})^{\uparrow x_i, s_i} \cap M = \emptyset$ by Claim 9. Moreover, we are given that $(y, t) \not\prec_{\mathcal{R}} (x_i, s_i)$, and thus we have not added **before**(y, t, x_i, s_i) to $\text{caus}_{\mathcal{R}}^{[i]}$, and by extension also not to M (since $\text{caus}_{\mathcal{R}}^{[i]}$ is the only part of M where we add **before**-facts with last two components x_i and s_i). Thus overall $V''(neg_{\varphi''}) \cap M = \emptyset$, as desired.
- Now we show $V''(pos_{\varphi''}) \subseteq N$. By construction of φ'' , the set $V''(pos_{\varphi''})$ consists of the facts $V(pos_{\varphi})^{\uparrow x_i, s_i}$, **all**(y) and **time**(t). First, we immediately have **time**(t) $\in \text{decl}(H) \subseteq N$. Also, by definition of $\text{msg}_{\mathcal{R}}^{[i]}$, y is a valid addressee and thus **all**(y) $\in \text{decl}(H) \subseteq N$. Finally, because V is satisfying for φ , we have $V(pos_{\varphi}) \subseteq D_i$. Thus $V(pos_{\varphi})^{\uparrow x_i, s_i} \subseteq D_i^{\uparrow x_i, s_i}$, and we are given that $D_i^{\uparrow x_i, s_i} \subseteq N$. Thus overall $V''(pos_{\varphi''}) \subseteq N$.

□

A.2 Inclusion $N \subseteq M$

In this section we show that $N \subseteq M$. By definition, $N = G(\text{decl}(H))$. Following the semantics of positive Datalog⁺ programs in Section 3.2.1, we can view N as a fixpoint, i.e., $N = \bigcup_{l \in \mathbb{N}} N_l$, where $N_0 = \text{decl}(H)$, and for each $l \geq 1$ the set N_l is obtained by applying the immediate consequence operator of G to N_{l-1} . This implies $N_{l-1} \subseteq N_l$ for each $l \geq 1$. We show by induction on $l = 0, 1, \dots$, that $N_l \subseteq M$. For the base case ($l = 0$), we immediately have $N_0 = \text{decl}(H) \subseteq M$. For the induction hypothesis, we assume for some $l \geq 1$ that $N_{l-1} \subseteq M$. For the inductive step, we show that $N_l \subseteq N$. Specifically, we divide the facts of $N_l \setminus N_{l-1}$ into groups based on their predicate, and for each group we show inclusion in M . As for terminology, we call a ground rule $\psi \in G$ *active* on N_{l-1} if $pos_{\psi} \subseteq N_{l-1}$. The numbered claims we will refer to can be found in Section A.2.5.

A.2.1 Causality

Let **before**(x, s, y, t) $\in N_l \setminus N_{l-1}$. It is sufficient to show that $(x, s) \prec_{\mathcal{R}} (y, t)$ because then **before**(x, s, y, t) $\in \text{caus}_{\mathcal{R}}^{[i]} \subseteq M$ where $i = \text{glob}_{\mathcal{R}}(y, t)$. We have the following cases:

Local edge The **before**-fact was derived by a ground rule in G of the form (7) (local edge). This implies $x = y$ and $t = s + 1$. Then $(x, s) \prec_{\mathcal{R}} (y, t)$ by definition of $\prec_{\mathcal{R}}$.

Message edge The **before**-fact was derived by a ground rule in G of the form (10) (message edge):

$$\text{before}(x, s, y, t) \leftarrow \text{chosen}_R(x, s, y, t, \bar{a}).$$

Since this rule is active on N_{l-1} , we have $\text{chosen}_R(x, s, y, t, \bar{a}) \in N_{l-1}$. By applying the induction hypothesis, we have $\text{chosen}_R(x, s, y, t, \bar{a}) \in M$. Denoting

$j = \text{glob}_{\mathcal{R}}(x, s)$, the set $\text{snd}_{\mathcal{R}}^{[j]}$ is the only part of M where we could have added this fact. This implies that x during its step s sends a message to y , and this message arrives at local step t of y . Then $(x, s) \prec_{\mathcal{R}} (y, t)$ by definition of $\prec_{\mathcal{R}}$.

Transitive edge The **before**-fact was derived by a ground rule in G of the form (8) (transitive edge):

$$\mathbf{before}(x, s, y, t) \leftarrow \mathbf{before}(x, s, z, u), \mathbf{before}(z, u, y, t).$$

Since this rule is active on N_{l-1} , its body facts are in N_{l-1} . By applying the induction hypothesis, we have $\mathbf{before}(x, s, z, u) \in M$ and $\mathbf{before}(z, u, y, t) \in M$. The only places we could have added these facts to M are in the sets $\text{caus}_{\mathcal{R}}^{[j]}$ and $\text{caus}_{\mathcal{R}}^{[k]}$ respectively, where $j = \text{glob}_{\mathcal{R}}(z, u)$ and $k = \text{glob}_{\mathcal{R}}(y, t)$. By construction of the sets $\text{caus}_{\mathcal{R}}^{[j]}$ and $\text{caus}_{\mathcal{R}}^{[k]}$ we respectively have that $(x, s) \prec_{\mathcal{R}} (z, u)$ and $(z, u) \prec_{\mathcal{R}} (y, t)$, and thus by transitivity $(x, s) \prec_{\mathcal{R}} (y, t)$, as desired.

A.2.2 Finite Messages

Senders Let $\mathbf{hasSender}(x, s, y, t) \in N_l \setminus N_{l-1}$. This fact can only have been derived by a ground rule in G of the form (11):

$$\mathbf{hasSender}(x, s, y, t) \leftarrow \mathbf{chosen}_R(y, t, x, s, \bar{a}).$$

Since this rule is active on N_{l-1} , we have $\mathbf{chosen}_R(y, t, x, s, \bar{a}) \in N_{l-1}$. By applying the induction hypothesis, we have $\mathbf{chosen}_R(y, t, x, s, \bar{a}) \in M$. We can only have added this fact in the set $\text{snd}_{\mathcal{R}}^{[i]}$ with $i = \text{glob}_{\mathcal{R}}(y, t)$. This means that y during its step t sends a message $R(\bar{a})$ to x , and this message arrives during step s of x . Hence, denoting $j = \text{glob}_{\mathcal{R}}(x, s)$, we have $(y, t) \in \text{senders}_{\mathcal{R}}^{[j]}$ (with $\text{senders}_{\mathcal{R}}^{[j]}$ as defined in Section 5.2.2). Thus we have added the fact $\mathbf{hasSender}(x, s, y, t) \in \text{fin}_{\mathcal{R}}^{[j]} \subseteq M$, as desired.

Comparison of timestamps Let $\mathbf{isSmaller}(x, s, y, t) \in N_l \setminus N_{l-1}$. This fact can only have been derived by a ground rule in G of the form (12):

$$\begin{aligned} \mathbf{isSmaller}(x, s, y, t) &\leftarrow \mathbf{hasSender}(x, s, y, t), \mathbf{hasSender}(x, s, y, t'), \\ &t < t'. \end{aligned}$$

Since this rule is active on N_{l-1} , its body facts are in N_{l-1} . By applying the induction hypothesis, we have $\mathbf{hasSender}(x, s, y, t) \in M$ and $\mathbf{hasSender}(x, s, y, t') \in M$. The only part of M where we could have added these facts is the set $\text{fin}_{\mathcal{R}}^{[i]}$ with $i = \text{glob}_{\mathcal{R}}(x, s)$. By construction of the set $\text{fin}_{\mathcal{R}}^{[i]}$, this implies that $(y, t) \in \text{senders}_{\mathcal{R}}^{[i]}$ and $(y, t') \in \text{senders}_{\mathcal{R}}^{[i]}$. Because $(t < t') \in N_{l-1}$, we more specifically know that $(t < t') \in \text{decl}(H)$, which implies $t < t'$. Thus we have added $\mathbf{isSmaller}(x, s, y, t) \in \text{fin}_{\mathcal{R}}^{[i]}$, as desired.

Maximum timestamp Let $\mathbf{hasMax}(x, s, y) \in N_l \setminus N_{l-1}$. This fact can only have been derived by a ground rule in G of the form (13):

$$\mathbf{hasMax}(x, s, y) \leftarrow \mathbf{hasSender}(x, s, y, t).$$

Since this rule is active on N_{l-1} , we have $\text{hasSender}(x, s, y, t) \in N_{l-1}$. By applying the induction hypothesis, we have $\text{hasSender}(x, s, y, t) \in M$. The only part of M where we could have added this fact, is the set $\text{fin}_{\mathcal{R}}^{[i]}$ with $i = \text{glob}_{\mathcal{R}}(x, s)$. Thus $(y, t) \in \text{senders}_{\mathcal{R}}^{[i]}$, and y is a sender-node mentioned in $\text{senders}_{\mathcal{R}}^{[i]}$. Hence, we have added $\text{hasMax}(x, s, y) \in \text{fin}_{\mathcal{R}}^{[i]} \subseteq M$, as desired.

Receive infinite Let $\text{rcvInf}(x, s) \in N_l \setminus N_{l-1}$. This fact can only have been derived by a ground rule in G of the form (14):

$$\text{rcvInf}(x, s) \leftarrow \text{hasSender}(x, s, y, t).$$

Since this rule is active on N_{l-1} , we have $\text{hasSender}(x, s, y, t) \in N_{l-1}$. By applying the induction hypothesis, we have $\text{hasSender}(x, s, y, t) \in M$. The only part of M where we could have added this fact, is the set $\text{fin}_{\mathcal{R}}^{[i]}$ with $i = \text{glob}_{\mathcal{R}}(x, s)$. Thus $(y, t) \in \text{senders}_{\mathcal{R}}^{[i]}$. Moreover, because the rule (14) contains a negative hasMax -atom in the body, and the above ground rule is in G , it must be that $\text{hasMax}(x, s, y) \notin M$, and thus $\text{hasMax}(x, s, y) \notin \text{fin}_{\mathcal{R}}^{[i]}$. But since y is a sender-node mentioned in $\text{senders}_{\mathcal{R}}^{[i]}$, the absence of $\text{hasMax}(x, s, y)$ from $\text{fin}_{\mathcal{R}}^{[i]}$ is impossible. Therefore this case can not occur.

A.2.3 Regular Facts

Let $R(x, s, \bar{a}) \in (N_l \setminus N_{l-1})|_{\text{sch}(\mathcal{P})^{\text{LT}}}$. The fact $R(x, s, \bar{a})$ has been derived by a ground rule $\psi \in G$ that is active on N_{l-1} . Because $\psi \in G$, there is a rule $\varphi \in \text{pure}(\mathcal{P})$ and valuation V such that ψ is obtained from φ by applying V and by subsequently removing the negative (ground) body atoms, and such that $V(\text{neg}_{\varphi}) \cap M = \emptyset$. We have the following cases:

Deductive Rule φ is of the form (1). Let $\varphi' \in \text{deduc}_{\mathcal{P}}$ be the original deductive rule corresponding to φ . By construction of φ out of φ' , we can apply valuation V to φ' as well. Denote $i = \text{glob}_{\mathcal{R}}(x, s)$. We will show now that V is satisfying for φ' during transition i , which causes $V(\text{head}_{\varphi'}) = R(\bar{a}) \in D_i$ to be derived, and we obtain as desired:

$$R(x, s, \bar{a}) \in D_i^{\uparrow x, s} = D_i^{\uparrow x_i, s_i} = \text{duc}_{\mathcal{R}}^{[i]} \subseteq M.$$

By definition of syntactic stratification, relations mentioned in $\text{pos}_{\varphi'}$ are never computed in a stratum higher than R , and relations mentioned in $\text{neg}_{\varphi'}$ are computed in a strictly lower stratum than R . Thus, it is sufficient to show that $V(\text{pos}_{\varphi'}) \subseteq D_i$ and $V(\text{neg}_{\varphi'}) \cap D_i = \emptyset$.

First we show $V(\text{pos}_{\varphi'}) \subseteq D_i$. Because φ is of the form (1), all facts in $V(\text{pos}_{\varphi})$ are over $\text{sch}(\mathcal{P})^{\text{LT}}$ and have location specifier x and timestamp s . Moreover, since ψ is active on N_{l-1} , we have $\text{pos}_{\psi} = V(\text{pos}_{\varphi}) \subseteq N_{l-1}$. By applying the induction hypothesis, we have $V(\text{pos}_{\varphi}) \subseteq M$, and thus $V(\text{pos}_{\varphi})^{\downarrow} \subseteq D_i$ by Claim 12. We thus obtain $V(\text{pos}_{\varphi'}) \subseteq D_i$ since $V(\text{pos}_{\varphi})^{\downarrow} = V(\text{pos}_{\varphi'})$.

Next we show $V(\text{neg}_{\varphi'}) \cap D_i = \emptyset$. Because φ is of the form (1), all facts in $V(\text{neg}_{\varphi})$ are over $\text{sch}(\mathcal{P})^{\text{LT}}$ and have location specifier x and timestamp s . Moreover, by

choice of φ and V , we have $V(\text{neg}_\varphi) \cap M = \emptyset$, and thus $V(\text{neg}_\varphi)^\Downarrow \cap D_i = \emptyset$ by Claim 13. We thus obtain $V(\text{neg}_{\varphi'}) \cap D_i = \emptyset$ since $V(\text{neg}_\varphi)^\Downarrow = V(\text{neg}_{\varphi'})$.

Inductive Rule φ is of the form (2). Let $\varphi' \in \text{induc}_\mathcal{P}$ be the rule corresponding to φ . First, ψ contains in its body a fact of the form $\text{tsucc}(r, s)$. Since ψ is active on N_{l-1} , we have $\text{tsucc}(r, s) \in N_{l-1}$ and more specifically, $\text{tsucc}(r, s) \in \text{decl}(H)$. This implies that $s = r + 1$. Denote $i = \text{glob}_\mathcal{R}(x, r)$ and $j = \text{glob}_\mathcal{R}(x, s)$. Since $s = r + 1$, there are no transitions of node x between i and j . By the relationship between φ and φ' , we can apply V to φ' , and we will now show that V is satisfying for φ' during transition i . This results in $V(\text{head}_{\varphi'}) = R(\bar{a}) \in \text{induc}_\mathcal{P}(D_i) \subseteq \text{st}_{i+1}(x)$, and since $\text{st}_{i+1}(x) = \text{st}_j(x) \subseteq D_j$, we obtain $R(x, s, \bar{a}) \in D_j^{\uparrow x, s} = \text{duc}_\mathcal{R}^{[j]} \subseteq M$, as desired.

First we show $V(\text{pos}_{\varphi'}) \subseteq D_i$. Denote $I = V(\text{pos}_{\varphi'})|_{\text{sch}(\mathcal{P})^{\text{LT}}}$, which allows us to exclude the extra tsucc -fact in the body. All facts in I have location specifier x and timestamp r . Because ψ is active on N_{l-1} , we have $I \subseteq \text{pos}_\psi \subseteq N_{l-1}$, and by applying the induction hypothesis, we have $I \subseteq M$. Thus $I^\Downarrow \subseteq D_i$ by Claim 12. Hence, $V(\text{pos}_{\varphi'}) = I^\Downarrow \subseteq D_i$.

Secondly, showing that $V(\text{neg}_{\varphi'}) \cap D_i = \emptyset$ is like in the previous case, where φ is deductive.

Delivery Rule φ is of the form (6). Then ψ concretely looks as follows, where $(y, t) \in \mathcal{N} \times \mathbb{N}$:

$$R(x, s, \bar{a}) \leftarrow \text{chosen}_R(y, t, x, s, \bar{a}).$$

Since ψ is active on N_{l-1} , we have $\text{chosen}_R(y, t, x, s, \bar{a}) \in N_{l-1}$, and by applying the induction hypothesis, we have $\text{chosen}_R(y, t, x, s, \bar{a}) \in M$. The only part of M where we could have added this fact, is $\text{snd}_\mathcal{R}^{[i]}$ with $i = \text{glob}_\mathcal{R}(y, t)$. This implies that x will receive $R(\bar{a})$ during its local step s , thus during transition $j = \text{glob}_\mathcal{R}(x, s)$. Then, by the operational semantics, we have $R(\bar{a}) \in \text{untag}(m_j) \subseteq D_j$. Hence, $R(x, s, \bar{a}) \in D_j^{\uparrow x, s} = \text{duc}_\mathcal{R}^{[j]} \subseteq M$.

A.2.4 Sending

For a transition i of \mathcal{R} , let D_i denote the output of subprogram $\text{deduc}_\mathcal{P}$ during transition i .

Candidates Let $\text{cand}_R(x, s, y, t, \bar{a}) \in N_l \setminus N_{l-1}$. The fact $\text{cand}_R(x, s, y, t, \bar{a})$ is derived by a ground rule $\psi \in G$ of the form (9) that is active on N_{l-1} . Because $\psi \in G$, there is a rule $\varphi \in \text{pure}(\mathcal{P})$ and a valuation V such that ψ is obtained from φ by applying valuation V and by subsequently removing the negative (ground) body atoms, and so that $V(\text{neg}_\varphi) \cap M = \emptyset$. Denote $i = \text{glob}_\mathcal{R}(x, s)$. It is sufficient to show that $R(y, \bar{a}) \in \text{msg}_\mathcal{R}^{[i]}$ and $(y, t) \not\prec_\mathcal{R} (x, s)$, because then $\text{cand}_R(x, s, y, t, \bar{a}) \in \text{snd}_\mathcal{R}^{[i]} \subseteq M$, as desired.

First, we show $(y, t) \not\prec_\mathcal{R} (x, s)$. Because there is a negative **before**-atom in φ , the

existence of ψ in G implies that $\text{before}(y, t, x, s) \notin M$. Hence, $\text{before}(y, t, x, s) \notin \text{caus}_{\mathcal{R}}^{[i]}$. Then by construction of $\text{caus}_{\mathcal{R}}^{[i]}$ we obtain $(y, t) \not\prec_{\mathcal{R}} (x, s)$.

Secondly, we show $R(y, \bar{a}) \in \text{msg}_{\mathcal{R}}^{[i]}$. Let $\varphi' \in \mathcal{P}$ be the original asynchronous rule on which φ is based. Let $\varphi'' \in \text{async}_{\mathcal{P}}$ be the rule corresponding to φ' . It follows from the constructions of φ out of φ' and φ'' out of φ' that valuation V can be applied to φ'' . Note, $V(\text{head}_{\varphi''}) = R(y, \bar{a})$. We show that V is satisfying for φ'' during transition i on D_i , which gives $R(y, \bar{a}) \in \text{async}_{\mathcal{P}}(D_i)$. Moreover, the body of ψ contains the fact $\text{all}(y) \in \text{decl}(H)$, and thus $y \in \mathcal{N}$, making y a valid addressee. Hence, $R(y, \bar{a}) \in \text{msg}_{\mathcal{R}}^{[i]}$, as desired.

We have to show $V(\text{pos}_{\varphi''}) \subseteq D_i$ and $V(\text{neg}_{\varphi''}) \cap D_i = \emptyset$. Abbreviate $I_1 = V(\text{pos}_{\varphi})|_{\text{sch}(\mathcal{P})^{\text{LT}}}$ and $I_2 = V(\text{neg}_{\varphi})|_{\text{sch}(\mathcal{P})^{\text{LT}}}$. Note, $I_1^{\downarrow} = V(\text{pos}_{\varphi''})$ and $I_2^{\downarrow} = V(\text{neg}_{\varphi''})$. All facts in $I_1 \cup I_2$ have location specifier x and timestamp s .

- Because ψ is active on N_{l-1} , we have $I_1 \subseteq \text{pos}_{\psi} \subseteq N_{l-1}$, and thus $I_1 \subseteq M$ by the induction hypothesis. Then $V(\text{pos}_{\varphi''}) = I_1^{\downarrow} \subseteq D_i$ by Claim 12.
- By choice of φ and V , we have $I_2 \cap M = \emptyset$. Then $I_2^{\downarrow} \cap D_i = \emptyset$ by Claim 13, giving $V(\text{neg}_{\varphi''}) \cap D_i = \emptyset$.

Chosen Let $\text{chosen}_R(x, s, y, t, \bar{a}) \in N_l \setminus N_{l-1}$. This fact is derived by a ground rule ψ in G of the form (4):

$$\text{chosen}_R(x, s, y, t, \bar{a}) \leftarrow \text{cand}_R(x, s, y, t, \bar{a}).$$

Denote $i = \text{glob}_{\mathcal{R}}(x, s)$. We show that $R(y, \bar{a}) \in \text{msg}_{\mathcal{R}}^{[i]}$ and that t is the actual arrival timestamp of this message at y . Then $\text{chosen}_R(x, s, y, t, \bar{a}) \in \text{snd}_{\mathcal{R}}^{[i]} \subseteq M$, as desired.

First, since ψ is active on N_{l-1} , we have $\text{cand}_R(x, s, y, t, \bar{a}) \in N_{l-1}$, and thus $\text{cand}_R(x, s, y, t, \bar{a}) \in M$ by the induction hypothesis. The set $\text{snd}_{\mathcal{R}}^{[i]}$ is the only part of M where we could have added this fact, which implies $R(y, \bar{a}) \in \text{msg}_{\mathcal{R}}^{[i]}$ and $(y, t) \prec_{\mathcal{R}} (x, s)$.

We are left to show that t is the actual arrival timestamp of the message. Because $\psi \in G$, there is a rule $\varphi \in \text{pure}(\mathcal{P})$ and valuation V such that ψ is obtained from φ by applying V and by subsequently removing the negative (ground) body atoms, and so that $V(\text{neg}_{\varphi}) \cap M = \emptyset$. Now, because rule φ contains a negative **other**_R-atom in its body, we have $\text{other}_R(x, s, y, t, \bar{a}) \notin M$ and thus $\text{other}_R(x, s, y, t, \bar{a}) \notin \text{snd}_{\mathcal{R}}^{[i]}$. Since $R(y, \bar{a}) \in \text{msg}_{\mathcal{R}}^{[i]}$ and $(y, t) \prec_{\mathcal{R}} (x, s)$ (see above), the absence of this **other**_R-fact from $\text{snd}_{\mathcal{R}}^{[i]}$ can only be explained by the following: $t = \text{loc}_{\mathcal{R}}(j)$ with $j = \alpha_{\mathcal{R}}(i, y, R(\bar{a}))$, as desired.

Other Let $\text{other}_R(x, s, y, t, \bar{a}) \in N_l \setminus N_{l-1}$. This fact is derived by a ground rule ψ of the form (5):

$$\begin{aligned} \text{other}_R(x, s, y, t, \bar{a}) \quad \leftarrow \quad & \text{cand}_R(x, s, y, t, \bar{a}), \text{chosen}_R(x, s, y, t', \bar{a}), \\ & t \neq t'. \end{aligned}$$

We have $\text{cand}_R(x, s, y, t, \bar{a}) \in N_{l-1}$ and $\text{chosen}_R(x, s, y, t', \bar{a}) \in N_{l-1}$ since ψ is active on N_{l-1} , and these facts are thus also in M by the induction hypothesis.

Denote $i = \text{glob}_{\mathcal{R}}(x, s)$. The only part of M where we could have added these **cand** _{R} - and **chosen** _{R} -facts to M , is the set $\text{snd}_{\mathcal{R}}^{[i]}$. First, $\text{cand}_R(x, s, y, t, \bar{a}) \in \text{snd}_{\mathcal{R}}^{[i]}$ implies that $R(y, \bar{a}) \in \text{msg}_{\mathcal{R}}^{[i]}$ and $(y, t) \not\prec_{\mathcal{R}} (x, s)$. Second, $\text{chosen}_R(x, s, y, t', \bar{a}) \in \text{snd}_{\mathcal{R}}^{[i]}$ implies that t' is the real arrival timestamp of the message $R(\bar{a})$ at y . Finally, since ψ is active, we have $(t \neq t') \in \text{decl}(H)$, and thus $t \neq t'$. Therefore we have added **other** _{R} (x, s, y, t, \bar{a}) to $\text{snd}_{\mathcal{R}}^{[i]} \subseteq M$, as desired.

A.2.5 Subclaims

Claim 12

Let I be a set of facts over $\text{sch}(\mathcal{P})^{\text{LT}}$, all having the same location specifier $x \in \mathcal{N}$ and timestamp $s \in \mathbb{N}$. Denote $i = \text{glob}_{\mathcal{R}}(x, s)$. If $I \subseteq M$ then $I^{\downarrow} \subseteq D_i$, where D_i denotes the output of subprogram $\text{deduc}_{\mathcal{P}}$ during transition i of \mathcal{R} .

Proof

The only part of M where we add facts over $\text{sch}(\mathcal{P})^{\text{LT}}$ with location specifier x and timestamp s is $\text{duc}_{\mathcal{R}}^{[i]}$. Hence $I \subseteq \text{duc}_{\mathcal{R}}^{[i]} = D_i^{\uparrow x, s}$ and thus $I^{\downarrow} \subseteq D_i$. \square

Claim 13

Let I be a set of facts over $\text{sch}(\mathcal{P})^{\text{LT}}$, all having the same location specifier $x \in \mathcal{N}$ and timestamp $s \in \mathbb{N}$. Denote $i = \text{glob}_{\mathcal{R}}(x, s)$. If $I \cap M = \emptyset$ then $I^{\downarrow} \cap D_i = \emptyset$, where D_i denotes the output of subprogram $\text{deduc}_{\mathcal{P}}$ during transition i of \mathcal{R} .

Proof

First, $I \cap M = \emptyset$ implies $I \cap \text{duc}_{\mathcal{R}}^{[i]} = \emptyset$ because $\text{duc}_{\mathcal{R}}^{[i]} \subseteq M$. And since $\text{duc}_{\mathcal{R}}^{[i]} = D_i^{\uparrow x, s}$, we have $I \cap D_i^{\uparrow x, s} = \emptyset$. Finally, since the facts in $I \cup D_i^{\uparrow x, s}$ all have the same location specifier x and timestamp s , we obtain $I^{\downarrow} \cap D_i = \emptyset$. \square

Appendix B Model to Run: Proof Details

Consider the definitions and notations from Section 5.3. In this section we show that \mathcal{R} is a run of \mathcal{P} on input H , and that $\text{trace}(\mathcal{R}) = M|_{\text{sch}(\mathcal{P})^{\text{LT}}}$. We do this in several parts, where each part is placed in its own subsection:

- in Section B.2 we show $\rho_0 = \text{start}(\mathcal{P}, H)$;
- in Section B.3 we show that every transition of \mathcal{R} is valid; and,
- in Section B.4 we show $\text{trace}(\mathcal{R}) = M|_{\text{sch}(\mathcal{P})^{\text{LT}}}$.

Before we start, the next subsection gives definitions and notations. The numbered claims we will refer to can be found in Section B.5.

B.1 Definitions and Notations

Using notations of Section 3.2.3, let G be the ground program $\text{ground}_M(C, I)$ where $C = \text{pure}(\mathcal{P})$ and $I = \text{decl}(H)$. By definition of M as a stable model, we have $M = G(I)$.

Let $\varphi \in \text{pure}(\mathcal{P})$ be a rule having its head atom over $\text{sch}(\mathcal{P})^{\text{LT}}$. From the construction of $\text{pure}(\mathcal{P})$, we know that φ belongs to exactly one of the following three cases:

- φ is of the form (1), i.e., *deductive*, recognizable as a rule in which only atoms over $\text{sch}(\mathcal{P})^{\text{LT}}$ are used, and in which the location and timestamp variable in the head are the same as in the body;
- φ is of the form (2), i.e., *inductive*, recognizable as a rule with a head atom over $\text{sch}(\mathcal{P})^{\text{LT}}$ and a **tsucc**-atom in the body;
- φ is of the form (6), i.e., a *delivery*, recognizable as a rule with a head atom over $\text{sch}(\mathcal{P})^{\text{LT}}$ and a **chosen_R**-fact in the body (with R the head-predicate).

The same classification of deductive, inductive and delivery rules can also be applied to the (positive) ground rules in G that have a ground head atom over $\text{sch}(\mathcal{P})^{\text{LT}}$.

Recall from the general remarks at the beginning of the appendix that we are working with a fixed (but arbitrary) syntactic stratification for the deductive rules. Stratum numbers start at 1. If $\varphi \in \text{pure}(\mathcal{P})$ is deductive, we can uniquely identify its stratum number as the stratum number of the original deductive rule in \mathcal{P} on which φ is based. Similarly, for deductive ground rules, we can also uniquely identify the stratum number as the stratum number of a corresponding non-ground rule in $\text{pure}(\mathcal{P})$.¹⁶

We call a ground rule $\psi \in G$ *active* if $\text{pos}_\psi \subseteq M$, which implies that $\text{head}_\psi \in M$ because M is stable. Now we define the following subsets of M :

- $M^{\text{duc},k}$: the head facts of all active deductive rules in G with stratum number less than or equal to k ;
- M^{ind} : the head facts of all active inductive rules in G ;
- M^{deliv} : the head facts of all active delivery rules in G .

This allows us to classify the facts in $M|_{\text{sch}(\mathcal{P})^{\text{LT}}}$ as being derived in a deductive manner, an inductive manner or being message deliveries. We also define:

$$M^\blacktriangle = M|_{\text{edb}(\mathcal{P})^{\text{LT}}} \cup M^{\text{ind}} \cup M^{\text{deliv}}.$$

For $(x, s) \in \mathcal{N} \times \mathbb{N}$, we write $I|^{x,s}$ to abbreviate $(I|_{\text{sch}(\mathcal{P})^{\text{LT}}})|^{x,s}$. So intuitively, when we select the facts with location specifier x and timestamp s , we are only interested in facts that provide these two components, which are the facts over $\text{sch}(\mathcal{P})^{\text{LT}}$.

Intuitively, for $i \in \mathbb{N}$, the set $(M^\blacktriangle)|^{x_i, s_i}$ is the input for the deductive rules during local step s_i of node x_i , consisting of (i) the *edb*-facts; (ii) the facts derived

¹⁶ We say a rather than *the* corresponding rule because there could be more than one. Indeed, multiple original deductive rules in $\text{pure}(\mathcal{P})$ could be mapped to the same positive ground rule after applying a valuation and removing their negative ground body atoms. But in any case, these non-ground rules will have the same head predicate. Hence, they have the same stratum.

by inductive rules during a previous step (if any) of x_i ; and, (iii) the delivered messages. The deductive rules then complete this information by deriving some new facts, that are visible within step s_i of x_i .

For a transition number i of \mathcal{R} , (i) we denote the source-configuration of transition i as $\rho_i = (st_i, bf_i)$; (ii) we denote the set of (tagged) messages delivered in transition i as m_i ; and, (iii) we denote $D_i = \text{deduc}_{\mathcal{P}}(st_i(x_i) \cup \text{untag}(m_i))$. For a number $k \in \mathbb{N}$, we write $D_i^{\rightarrow k}$ to denote the set of facts obtained by adding to $st_i(x_i) \cup \text{untag}(m_i)$ all facts derived in stratum 1 up to stratum k during the computation of D_i . To mirror this notation, we write $M^{\rightarrow k}$ to denote the set $M^{\blacktriangle} \cup M^{\text{duc}, k}$. For uniformity in the proofs, we will consider the case $k = 0$, which is an invalid stratum number, and this gives $D_i^{\rightarrow 0} = st_i(x_i) \cup \text{untag}(m_i)$ and $M^{\rightarrow 0} = M^{\blacktriangle}$.

B.2 Valid Start

We show that $\rho_0 = \text{start}(\mathcal{P}, H)$. Denote $\rho_0 = (st_0, bf_0)$. Let $x \in \mathcal{N}$. First we show $st_0(x) = H(x)$. By definition,

$$st_0(x) = ((M|_{\text{edb}(\mathcal{P})^{\text{LT}}})|^{x,s} \cup M^{\text{ind}}|^{x,s})^{\Downarrow}$$

with $s = \text{loc}_M(0, x)$. Note, $s = 0$ because no elements of $\mathcal{N} \times \mathbb{N}$ with first component x have an ordinal strictly less than 0 in the total order $<_M$. Now, there can be no ground inductive rules in G that derive facts with head timestamp 0 because it follows from the construction of $\text{decl}(H)$ that the second component of a tsucc -fact is always strictly larger than 0. Therefore $M^{\text{ind}}|^{x,s} = \emptyset$, and thus $st_0(x) = ((M|_{\text{edb}(\mathcal{P})^{\text{LT}}})|^{x,s})^{\Downarrow}$. Then by Claim 14 we have $st_0(x) = (H(x)^{\uparrow x,s})^{\Downarrow} = H(x)$, as desired.

Now we show $bf_0(x) = \emptyset$. By definition, $bf_0(x)$ is

$$\{(glob_M(y, t), R(\bar{a})) \mid \exists u : \text{chosen}_R(y, t, x, u, \bar{a}) \in M, \\ glob_M(y, t) < 0 \leq glob_M(x, u)\}.$$

By definition of function $glob_M(\cdot)$, all facts of the form $\text{chosen}_R(y, t, x, u, \bar{a}) \in M$ satisfy $glob_M(y, t) \geq 0$. Hence, $bf_0(x) = \emptyset$.

We conclude that $\rho_0 = \text{start}(\mathcal{P}, H)$.

B.3 Valid Transition

Let $i \in \mathbb{N}$. We show that $(\rho_i, x_i, m_i, i, \rho_{i+1})$ is a valid transition. Denote $\rho_i = (st_i, bf_i)$ and $\rho_{i+1} = (st_{i+1}, bf_{i+1})$.

We start by showing $m_i \subseteq bf_i(x_i)$. Let $(j, \mathbf{f}) \in m_i$. By definition of m_i , there is a fact of the form $\text{chosen}_R(y, t, z, u, \bar{a}) \in M$ with $glob_M(z, u) = i$ such that $j = glob_M(y, t)$ and $\mathbf{f} = R(\bar{a})$. Note, $glob_M(z, u) = i$ implies $z = x_i$ and $u = s_i$. Now, because rules in $\text{pure}(\mathcal{P})$ of the form (10) are always positive, the following ground rule is in G , which is of the form (10):

$$\text{before}(y, t, x_i, s_i) \leftarrow \text{chosen}_R(y, t, x_i, s_i, \bar{a}).$$

Since its body is in M , this rule derives $\text{before}(y, t, x_i, s_i) \in M$. Hence $(y, t) \prec_M$

(x_i, s_i) by definition of \prec_M . Moreover, $<_M$ respects \prec_M , and thus $(y, t) <_M (x_i, s_i)$, which implies $glob_M(y, t) < glob_M(x_i, s_i)$. And since $glob_M(x_i, s_i) = i$, we overall have

$$glob_M(y, t) < i \leq glob_M(x_i, s_i).$$

Therefore $(j, \mathbf{f}) \in bf_i(x_i)$.

Now, because $m_i \subseteq bf_i(x_i)$, and because transitions are deterministic once the active node and delivered messages are fixed, we can consider the unique result configuration $\rho = (st, bf)$ such that $(\rho_i, x_i, m_i, i, \rho)$ is a valid transition. We are left to show $\rho_{i+1} = \rho$. We divide the work in two parts: for each $x \in \mathcal{N}$, we show that (i) $st_{i+1}(x) = st(x)$, and (ii) $bf_{i+1}(x) = bf(x)$.

B.3.1 State

Let $x \in \mathcal{N}$. We show $st_{i+1}(x) = st(x)$. Denote $s = loc_M(i+1, x)$. By definition,

$$st_{i+1}(x) = ((M|_{edb(\mathcal{P})^{LT}})^{x,s} \cup M^{\text{ind}|x,s})^{\Downarrow}.$$

Case $x \neq x_i$. By definition, $st(x) = st_i(x)$. Hence, it suffices to show $st_{i+1}(x) = st_i(x)$. Since $x \neq x_i$, the number of pairs from $\mathcal{N} \times \mathbb{N}$ containing node x that come strictly before ordinal $i+1$ is the same as the number of pairs containing node x that come strictly before ordinal i . Formally: $s = loc_M(i+1, x) = loc_M(i, x)$. Thus the right-hand side in the previous equation equals $st_i(x)$, and the result is obtained.

Case $x = x_i$. By definition, $st(x) = H(x) \cup \text{induc}_{\mathcal{P}}(D_i)$. Referring to the definition of $st_{i+1}(x)$ from above, by Claim 14 we have

$$(M|_{edb(\mathcal{P})^{LT}})^{x,s} = H(x)^{\uparrow x,s}.$$

If we can also show $M^{\text{ind}|x,s} = \text{induc}_{\mathcal{P}}(D_i)^{\uparrow x,s}$, then we overall have, as desired:

$$\begin{aligned} st_{i+1}(x) &= ((M|_{edb(\mathcal{P})^{LT}})^{x,s} \cup M^{\text{ind}|x,s})^{\Downarrow} \\ &= H(x) \cup \text{induc}_{\mathcal{P}}(D_i) \\ &= st(x). \end{aligned}$$

Since $x = x_i$, we have $s = loc_M(i+1, x_i) = loc_M(i, x_i) + 1$, and using that $loc_M(i, x_i) = s_i$ (Claim 15), we have $s = s_i + 1$. Now, Claim 16 and Claim 19 together show $M^{\text{ind}|x_i, s_i+1} = \text{induc}_{\mathcal{P}}(D_i)^{\uparrow x_i, s_i+1}$.

B.3.2 Buffer

Let $x \in \mathcal{N}$. We show $bf_{i+1}(x) = bf(x)$. Denote

$$\delta^{i \rightarrow x} = \{(i, R(\bar{a})) \mid R(x, \bar{a}) \in \text{async}_{\mathcal{P}}(D_i)\}.$$

Like in the operational semantics, $\delta^{i \rightarrow x}$ denotes the (tagged) messages that are sent to x during transition i .

Case $x \neq x_i$. By definition, $bf(x) = bf_i(x) \cup \delta^{i \rightarrow x}$. We start by showing $bf(x) \subseteq bf_{i+1}(x)$. Let $(j, \mathbf{f}) \in bf(x)$. Denote $\mathbf{f} = R(\bar{a})$.

- Suppose $(j, \mathbf{f}) \in bf_i(x)$. By definition of $bf_i(x)$, there are values $y \in \mathcal{N}$, $t \in \mathbb{N}$ and $u \in \mathbb{N}$ such that $\text{chosen}_R(y, t, x, u, \bar{a}) \in M$ and $j = \text{glob}_M(y, t) < i \leq \text{glob}_M(x, u)$. Now, since $x \neq x_i$, we more specifically have $i < \text{glob}_M(x, u)$ and thus $i + 1 \leq \text{glob}_M(x, u)$. Therefore $(j, \mathbf{f}) \in bf_{i+1}(x)$, as desired.
- Suppose $(j, \mathbf{f}) \in \delta^{i \rightarrow x}$. By definition of $\delta^{i \rightarrow x}$, this implies $j = i$ and $R(x, \bar{a}) \in \text{async}_{\mathcal{P}}(D_i)$. Then $(j, \mathbf{f}) = (i, R(\bar{a})) \in bf_{i+1}(x)$ by Claim 20, as desired.

Secondly, we show $bf_{i+1}(x) \subseteq bf(x)$. Let $(j, \mathbf{f}) \in bf_{i+1}(x)$. Denote $\mathbf{f} = R(\bar{a})$. By definition of $bf_{i+1}(x)$, there are values $y \in \mathcal{N}$, $t \in \mathbb{N}$ and $u \in \mathbb{N}$ such that $\text{chosen}_R(y, t, x, u, \bar{a}) \in M$ and $j = \text{glob}_M(y, t) < i + 1 \leq \text{glob}_M(x, u)$. So $j \leq i$. We have the following cases:

- Suppose $j < i$. Thus $\text{glob}_M(y, t) < i$. This immediately gives $(j, \mathbf{f}) \in bf_i(x) \subseteq bf(x)$, as desired.
- Suppose $j = i$. Then $R(x, \bar{a}) \in \text{async}_{\mathcal{P}}(D_i)$ by Claim 21. This implies that $(j, \mathbf{f}) = (i, R(\bar{a})) \in \delta^{i \rightarrow x} \subseteq bf(x)$, as desired.

Case $x = x_i$. By definition, $bf(x) = (bf_i(x) \setminus m_i) \cup \delta^{i \rightarrow x}$. Some parts of the reasoning are similar to the case $x \neq x_i$. We refer to shared subclaims where possible.

We start by showing $bf(x) \subseteq bf_{i+1}(x)$. Let $(j, \mathbf{f}) \in bf(x)$. Denote $\mathbf{f} = R(\bar{a})$. We have the following cases:

- Suppose $(j, \mathbf{f}) \in bf_i(x) \setminus m_i$. Thus $(j, \mathbf{f}) \in bf_i(x)$ and $(j, \mathbf{f}) \notin m_i$. Here, $(j, \mathbf{f}) \in bf_i(x)$ implies there are values $y \in \mathcal{N}$, $t \in \mathbb{N}$ and $u \in \mathbb{N}$ such that $\text{chosen}_R(y, t, x, u, \bar{a}) \in M$ and $j = \text{glob}_M(y, t) < i \leq \text{glob}_M(x, u)$. Also, $(j, \mathbf{f}) \notin m_i$ implies $\text{glob}_M(x, u) \neq i$. Hence, $i + 1 \leq \text{glob}_M(x, u)$ and we obtain $(j, \mathbf{f}) \in bf_{i+1}(x)$, as desired.
- Suppose $(j, \mathbf{f}) \in \delta^{i \rightarrow x}$. By definition of $\delta^{i \rightarrow x}$, we have $j = i$ and $R(x, \bar{a}) \in \text{async}_{\mathcal{P}}(D_i)$. By Claim 20 we then have $(i, R(\bar{a})) \in bf_{i+1}(x)$, as desired.

Secondly, we show $bf_{i+1}(x) \subseteq bf(x)$. Let $(j, \mathbf{f}) \in bf_{i+1}(x)$. Denote $\mathbf{f} = R(\bar{a})$. By definition of $bf_{i+1}(x)$, there are values $y \in \mathcal{N}$, $t \in \mathbb{N}$ and $u \in \mathbb{N}$ such that $\text{chosen}_R(y, t, x, u, \bar{a}) \in M$ and $j = \text{glob}_M(y, t) < i + 1 \leq \text{glob}_M(x, u)$. Now we look at the cases for j :

- Suppose $j < i$. This gives us $\text{glob}_M(y, t) < i \leq \text{glob}_M(x, u)$, which implies $(j, \mathbf{f}) \in bf_i(x)$. Moreover, $i + 1 \leq \text{glob}_M(x, u)$ gives $\text{glob}_M(x, u) \neq i$. Hence, $(j, \mathbf{f}) \notin m_i$. Taken together, we now have $(j, \mathbf{f}) \in bf_i(x) \setminus m_i \subseteq bf(x)$.
- Suppose $j = i$. Then $(i, R(\bar{a})) \in bf_{i+1}(x)$, and by Claim 21 we obtain that $R(x, \bar{a}) \in \text{async}_{\mathcal{P}}(D_i)$. Therefore $(j, \mathbf{f}) = (i, R(\bar{a})) \in \delta^{i \rightarrow x} \subseteq bf(x)$, as desired.

B.4 Trace

In this section we show $\text{trace}(\mathcal{R}) = M|_{\text{sch}(\mathcal{P})^{\text{LT}}}$. Recall from Section 5.1.5 that

$$\text{trace}(\mathcal{R}) = \bigcup_{i \in \mathbb{N}} (D_i)^{\uparrow x_i, \text{loc}_{\mathcal{R}}(i)}.$$

For each $i \in \mathbb{N}$, $\text{loc}_{\mathcal{R}}(i)$ is the number of transitions in \mathcal{R} before i in which x_i is also the active node. From the construction of \mathcal{R} we know $\text{loc}_{\mathcal{R}}(i) = \text{loc}_M(i, x_i)$; indeed, $\text{loc}_M(i, x_i)$ counts the number of pairs in $\mathcal{N} \times \mathbb{N}$ with node x_i that have an ordinal strictly smaller than i , which is precisely the number of transitions in \mathcal{R} with active node x_i that come before i . Moreover, by Claim 15 we have $\text{loc}_M(i, x_i) = s_i$. Hence,

$$\text{trace}(\mathcal{R}) = \bigcup_{i \in \mathbb{N}} (D_i)^{\uparrow x_i, s_i}.$$

Thus, by Claim 22:

$$\text{trace}(\mathcal{R}) = \bigcup_{i \in \mathbb{N}} M|^{x_i, s_i}.$$

For the next step, let us denote $A = \{(x_i, s_i) \mid i \in \mathbb{N}\}$. We show $A = \mathcal{N} \times \mathbb{N}$. First, we have $A \subseteq \mathcal{N} \times \mathbb{N}$ because $x_i \in \mathcal{N}$ and $s_i \in \mathbb{N}$ for each $i \in \mathbb{N}$. Now, let $(x, s) \in \mathcal{N} \times \mathbb{N}$. Denote $i = \text{glob}_M(x, s)$. By definition, $x_i = x$ and $s_i = s$. Hence $(x, s) = (x_i, s_i) \in A$. Now we may write:

$$\begin{aligned} \text{trace}(\mathcal{R}) &= \bigcup_{(x, s) \in A} M|^{x, s} \\ &= \bigcup_{(x, s) \in \mathcal{N} \times \mathbb{N}} M|^{x, s}. \end{aligned}$$

Finally, because M is well-formed (see Section 5.3), for each $R(v, w, \bar{a}) \in M|_{\text{sch}(\mathcal{P})^{\text{LT}}}$ we have $v \in \mathcal{N}$ and $w \in \mathbb{N}$. We obtain, as desired:

$$\text{trace}(\mathcal{R}) = M|_{\text{sch}(\mathcal{P})^{\text{LT}}}.$$

B.5 Subclaims

Claim 14

Let $x \in \mathcal{N}$ and $s \in \mathbb{N}$. We have $(M|_{\text{edb}(\mathcal{P})^{\text{LT}}})|^{x, s} = H(x)^{\uparrow x, s}$.

Proof

First, by construction of $\text{decl}(H)$ we have $(\text{decl}(H)|_{\text{edb}(\mathcal{P})^{\text{LT}}})|^{x, s} = H(x)^{\uparrow x, s}$. Because $\text{decl}(H) \subseteq M$, and because facts over $\text{edb}(\mathcal{P})^{\text{LT}}$ can not be derived by rules in $\text{pure}(\mathcal{P})$, we have $M|_{\text{edb}(\mathcal{P})^{\text{LT}}} = \text{decl}(H)|_{\text{edb}(\mathcal{P})^{\text{LT}}}$. Hence,

$$(M|_{\text{edb}(\mathcal{P})^{\text{LT}}})|^{x, s} = (\text{decl}(H)|_{\text{edb}(\mathcal{P})^{\text{LT}}})|^{x, s} = H(x)^{\uparrow x, s}.$$

□

Claim 15

Let $i \in \mathbb{N}$. We have $s_i = \text{loc}_M(i, x_i)$.

Proof

Recall that $(x_i, s_i) \in \mathcal{N} \times \mathbb{N}$ is the unique pair at ordinal i in $<_M$, i.e., $\text{glob}_M(x_i, s_i) = i$. Suppose we would know for all $s \in \mathbb{N}$ and $t \in \mathbb{N}$ that $s < t$ implies $\text{glob}_M(x_i, s) < \text{glob}_M(x_i, t)$. Then $\text{loc}_M(i, x_i)$, which is

$$|\{s \in \mathbb{N} \mid \text{glob}_M(x, s) < i\}|,$$

is precisely

$$|\{s \in \mathbb{N} \mid s < s_i\}|.$$

The latter is just s_i .

We are left to show for any $s \in \mathbb{N}$ and $t \in \mathbb{N}$ that $s < t$ implies $\text{glob}_M(x_i, s) < \text{glob}_M(x_i, t)$. It is actually sufficient to show for any $s \in \mathbb{N}$ that $(x_i, s) \prec_M (x_i, s+1)$. Indeed, this would imply for any $t \in \mathbb{N}$ with $s < t$ that

$$(x_i, s) \prec_M (x_i, s+1) \prec_M (x_i, s+2) \prec_M \dots \prec_M (x_i, t).$$

And since \prec_M is a partial order, it is transitive, and thus $(x_i, s) \prec_M (x_i, t)$. Next, since $<_M$ respects \prec_M , we obtain $(x_i, s) <_M (x_i, t)$ and thus $\text{glob}_M(x_i, s) < \text{glob}_M(x_i, t)$, as desired. To show $(x_i, s) \prec_M (x_i, s+1)$, we observe that the rule (7) in $\text{pure}(\mathcal{P})$ is positive. Hence, for any $s \in \mathbb{N}$, the following ground rule is always in G , and it derives $\text{before}(x_i, s, x_i, s+1) \in M$ because $\text{all}(x_i) \in \text{decl}(H)$ and $\text{tsucc}(s, s+1) \in \text{decl}(H)$:

$$\text{before}(x_i, s, x_i, s+1) \leftarrow \text{all}(x_i), \text{tsucc}(s, s+1).$$

Thus $(x_i, s) \prec_M (x_i, s+1)$ by definition of \prec_M . \square

Claim 16

Let $i \in \mathbb{N}$. We have $M^{\text{ind}}|_{x_i, s_i+1} \subseteq \text{induc}_{\mathcal{P}}(D_i)^{\uparrow x_i, s_i+1}$.

Proof

Let $\mathbf{f} \in M^{\text{ind}}|_{x_i, s_i+1}$. We show $\mathbf{f} \in \text{induc}_{\mathcal{P}}(D_i)^{\uparrow x_i, s_i+1}$.

By definition of M^{ind} , there is an active *inductive* ground rule $\psi \in G$ with $\text{head}_{\psi} = \mathbf{f}$. Because $\psi \in G$, there is a rule $\varphi \in \text{pure}(\mathcal{P})$ and a valuation V so that ψ can be obtained from φ by applying V and by subsequently removing all negative (ground) body literals, and so that $V(\text{neg}_{\varphi}) \cap M = \emptyset$. The rule φ must be of the form (2), which implies that V must assign x_i and s_i to the body location and timestamp variable respectively, and that it must assign x_i and $s_i + 1$ to the head location and timestamp variable respectively.

Let $\varphi' \in \mathcal{P}$ be the original inductive rule on which φ is based. Let $\varphi'' \in \text{induc}_{\mathcal{P}}$ be the rule corresponding to φ' . It follows from the construction of φ out of φ' and φ'' out of φ' that valuation V can also be applied to rule φ'' . Indeed, rule φ just has more variables for the location and timestamps. We show that V is satisfying

for φ'' with respect to D_i , so that φ'' and V together derive $V(\text{head}_{\varphi''}) = \mathbf{f}^\downarrow \in \text{induc}_{\mathcal{P}}(D_i)$, which gives $\mathbf{f} \in \text{induc}_{\mathcal{P}}(D_i)^{\uparrow x_i, s_i+1}$, as desired.

We must concretely show $V(\text{pos}_{\varphi''}) \subseteq D_i$ and $V(\text{neg}_{\varphi''}) \cap D_i = \emptyset$. We start by showing $V(\text{pos}_{\varphi''}) \subseteq D_i$. From the relationship between ψ , φ and φ'' , we know that

$$\text{pos}_{\psi}|_{\text{sch}(\mathcal{P})^{\text{LT}}} = V(\text{pos}_{\varphi})|_{\text{sch}(\mathcal{P})^{\text{LT}}} = V(\text{pos}_{\varphi''})^{\uparrow x_i, s_i}.$$

Since ψ is active with respect to M , we have $\text{pos}_{\psi} \subseteq M$, and thus $V(\text{pos}_{\varphi''})^{\uparrow x_i, s_i} \subseteq M$. Then by Claim 17 we have $V(\text{pos}_{\varphi''}) \subseteq D_i$, as desired.

Now we show that $V(\text{neg}_{\varphi''}) \cap D_i = \emptyset$. By the relationship of φ and φ'' , we have $V(\text{neg}_{\varphi''})^{\uparrow x_i, s_i} = V(\text{neg}_{\varphi})$. By choice of φ and V , we have $V(\text{neg}_{\varphi}) \cap M = \emptyset$. Hence, $V(\text{neg}_{\varphi''})^{\uparrow x_i, s_i} \cap M = \emptyset$. Finally, by Claim 18, we have $V(\text{neg}_{\varphi''}) \cap D_i = \emptyset$, as desired. \square

Claim 17

Let $i \in \mathbb{N}$. Let I be a set of facts over $\text{sch}(\mathcal{P})^{\text{LT}}$ that all have location specifier x_i and timestamp s_i . If $I \subseteq M$ then $I^\downarrow \subseteq D_i$, with D_i as defined in Section B.1.

Proof

We are given $I \subseteq M$. By the assumptions on I , we more specifically have $I \subseteq M|^{x_i, s_i}$. Then by Claim 22 we have $I \subseteq (D_i)^{\uparrow x_i, s_i}$. Hence $I^\downarrow \subseteq D_i$, as desired. \square

Claim 18

Let $i \in \mathbb{N}$. Let I be a set of facts over $\text{sch}(\mathcal{P})^{\text{LT}}$ that all have location specifier x_i and timestamp s_i . If $I \cap M = \emptyset$ then $I^\downarrow \cap D_i = \emptyset$, with D_i as defined in Section B.1.

Proof

We are given that $I \cap M = \emptyset$. This implies $I \cap M|^{x_i, s_i} = \emptyset$. By Claim 22 we have $I \cap (D_i)^{\uparrow x_i, s_i} = \emptyset$. Hence, by the assumptions on I , we have $I^\downarrow \cap D_i = \emptyset$, as desired. \square

Claim 19

Let $i \in \mathbb{N}$. We have $\text{induc}_{\mathcal{P}}(D_i)^{\uparrow x_i, s_i+1} \subseteq M^{\text{ind}|x_i, s_i+1}$.

Proof

Let $\mathbf{f} \in \text{induc}_{\mathcal{P}}(D_i)$. We show that $\mathbf{f}^{\uparrow x_i, s_i+1} \in M^{\text{ind}|x_i, s_i+1}$.

Recall the semantics for $\text{induc}_{\mathcal{P}}$ from Section 5.1.2. Let $\varphi \in \text{induc}_{\mathcal{P}}$ and V be the rule and valuation that together derived $\mathbf{f} \in \text{induc}_{\mathcal{P}}(D_i)$. Let $\varphi' \in \mathcal{P}$ be the original inductive rule on which φ is based. Let $\varphi'' \in \text{pure}(\mathcal{P})$ be the inductive rule that in turn is based on φ' , which is of the form (2). Let V'' be the valuation for φ'' that is obtained by extending V to assign x_i and s_i to respectively the location and timestamp variables in the body, and to assign $s_i + 1$ to the head timestamp variable. Let ψ be the positive ground rule obtained from φ'' by applying the

valuation V'' , and by subsequently removing the negative (ground) body literals. Note that $head_\psi = V(head_\varphi)^{\uparrow x_i, s_i+1} = f^{\uparrow x_i, s_i+1}$. We will show that $\psi \in G$ and that $pos_\psi \subseteq M$, so that this ground rule derives $f^{\uparrow x_i, s_i+1} \in M$. And since ψ is inductive, we more specifically have $f^{\uparrow x_i, s_i+1} \in M^{\text{ind}}|_{x_i, s_i+1}$, as desired.

- For $\psi \in G$, we require $V''(neg_{\varphi''}) \cap M = \emptyset$. From the construction of rule φ'' , we have $V''(neg_{\varphi''}) = V(neg_\varphi)^{\uparrow x_i, s_i}$. We show $V(neg_\varphi)^{\uparrow x_i, s_i} \cap M = \emptyset$. Because V is satisfying for φ with respect to D_i , we have $V(neg_\varphi) \cap D_i = \emptyset$. This gives $V(neg_\varphi)^{\uparrow x_i, s_i} \cap (D_i)^{\uparrow x_i, s_i} = \emptyset$. Then $V(neg_\varphi)^{\uparrow x_i, s_i} \cap M|_{x_i, s_i} = \emptyset$ by Claim 22. Next, we obtain $V(neg_\varphi)^{\uparrow x_i, s_i} \cap M = \emptyset$ since $V(neg_\varphi)^{\uparrow x_i, s_i}$ contains only facts over $sch(\mathcal{P})^{\text{LT}}$ with location specifier x_i and timestamp s_i .
- Now we show $pos_\psi \subseteq M$. From the construction of rule φ'' , we have

$$pos_\psi = V''(pos_{\varphi''}) = V(pos_\varphi)^{\uparrow x_i, s_i} \cup \{\text{tsucc}(s_i, s_i + 1)\}.$$

We immediately have $\text{tsucc}(s_i, s_i + 1) \in \text{decl}(H) \subseteq M$. Moreover, since V is satisfying for φ with respect to D_i , we have $V(pos_\varphi) \subseteq D_i$. Hence $V(pos_\varphi)^{\uparrow x_i, s_i} \subseteq (D_i)^{\uparrow x_i, s_i}$. By Claim 22 we then have $V(pos_\varphi)^{\uparrow x_i, s_i} \subseteq M|_{x_i, s_i} \subseteq M$, as desired.

□

Claim 20

Let $i \in \mathbb{N}$. Let $x \in \mathcal{N}$. For each $R(x, \bar{a}) \in \text{async}_{\mathcal{P}}(D_i)$, we have $(i, R(\bar{a})) \in \text{bf}_{i+1}(x)$.

Proof

The main approach of this proof is as follows. We will show there is a timestamp $u \in \mathbb{N}$ such that $\text{chosen}_R(x_i, s_i, x, u, \bar{a}) \in M$. Next, because rules of the form (10) are positive, in G there is always the following ground rule:

$$\text{before}(x_i, s_i, x, u) \leftarrow \text{chosen}_R(x_i, s_i, x, u, \bar{a}).$$

Thus if $\text{chosen}_R(x_i, s_i, x, u, \bar{a}) \in M$ then $\text{before}(x_i, s_i, x, u) \in M$, which implies $(x_i, s_i) \prec_M (x, u)$ by definition of \prec_M . Since $<_M$ respects \prec_M , we obtain $(x_i, s_i) <_M (x, u)$ and thus $\text{glob}_M(x_i, s_i) < \text{glob}_M(x, u)$. Also, since $\text{glob}_M(x_i, s_i) = i$, we overall get

$$\text{glob}_M(x_i, s_i) < i + 1 \leq \text{glob}_M(x, u),$$

which together with $\text{chosen}_R(x_i, s_i, x, u, \bar{a}) \in M$ gives $(\text{glob}_M(x_i, s_i), R(\bar{a})) = (i, R(\bar{a})) \in \text{bf}_{i+1}(x)$, as desired.

Now we are left to show that such a timestamp u exists. Recall the semantics for $\text{async}_{\mathcal{P}}$ from Section 5.1.2. Let $\varphi \in \text{async}_{\mathcal{P}}$ and V be a rule and valuation that together have derived $R(x, \bar{a}) \in \text{async}_{\mathcal{P}}(D_i)$. Let $\varphi' \in \mathcal{P}$ be the original asynchronous rule on which φ is based. Let $\varphi'' \in \text{pure}(\mathcal{P})$ be the rule obtained by applying transformation (9) to φ' . To continue, because \prec_M is well-founded, there are only a finite number of timestamps $v \in \mathbb{N}$ of node x such that $(x, v) \prec_M (x_i, s_i)$. So, there exists a timestamp $u \in \mathbb{N}$ such that $(x, u) \not\prec_M (x_i, s_i)$. Now, let V'' be the valuation for φ'' that is the extension of valuation V to assign x_i and s_i to the

body location variable and timestamp variable respectively (both belonging to the sender), and to assign u to the addressee arrival timestamp. Note that from the construction of φ'' we also know that V (and thus V'') assigns the value x to the addressee location variable and the tuple \bar{a} to the message contents. Let ψ denote the ground rule obtained by applying V'' to φ'' , and by subsequently removing the negative (ground) body literals. We will first show that $\psi \in G$, and then we show that $\text{pos}_\psi \subseteq M$, meaning that ψ derives $\text{head}_\psi = \text{cand}_R(x_i, s_i, x, u, \bar{a}) \in M$. Then Claim 24 can be applied to know that there is a timestamp u' , with possibly $u' = u$, such that $\text{chosen}_R(x_i, s_i, x, u', \bar{a}) \in M$, as desired.

In order for ψ to be in G , we require $V''(\text{neg}_{\varphi''}) \cap M = \emptyset$. It follows from the construction of φ'' out of φ' and φ out of φ' that

$$V''(\text{neg}_{\varphi''}) = V(\text{neg}_\varphi)^{\uparrow x_i, s_i} \cup \{\text{before}(x, u, x_i, s_i)\}.$$

We have $\text{before}(x, u, x_i, s_i) \notin M$ because $(x, u) \not\prec_M (x_i, s_i)$ by choice of u . Next, we show that $V(\text{neg}_\varphi)^{\uparrow x_i, s_i} \cap M = \emptyset$. Because V is satisfying for φ with respect to D_i , we have $V(\text{neg}_\varphi) \cap D_i = \emptyset$, and thus

$$V(\text{neg}_\varphi)^{\uparrow x_i, s_i} \cap (D_i)^{\uparrow x_i, s_i} = \emptyset.$$

Then, by Claim 22,

$$V(\text{neg}_\varphi)^{\uparrow x_i, s_i} \cap M|^{x_i, s_i} = \emptyset.$$

Since $V(\text{neg}_\varphi)^{\uparrow x_i, s_i}$ contains only facts over $\text{sch}(\mathcal{P})^{\text{LT}}$ with location specifier x_i and timestamp s_i , we have

$$V(\text{neg}_\varphi)^{\uparrow x_i, s_i} \cap M = \emptyset.$$

We now show $\text{pos}_\psi \subseteq M$. Note, $\text{pos}_\psi = V''(\text{pos}_{\varphi''})$. From the construction of φ'' we have

$$V''(\text{pos}_{\varphi''}) = V(\text{pos}_\varphi)^{\uparrow x_i, s_i} \cup \{\text{all}(x), \text{time}(u)\}.$$

Because $x \in \mathcal{N}$ and $u \in \mathbb{N}$, we immediately have $\{\text{all}(x), \text{time}(u)\} \subseteq \text{decl}(H) \subseteq M$. We are left to show $V(\text{pos}_\varphi)^{\uparrow x_i, s_i} \subseteq M$. Because V is satisfying for φ with respect to D_i , we have $V(\text{pos}_\varphi) \subseteq D_i$. Hence $V(\text{pos}_\varphi)^{\uparrow x_i, s_i} \subseteq (D_i)^{\uparrow x_i, s_i}$. By again using Claim 22 we then obtain $V(\text{pos}_\varphi)^{\uparrow x_i, s_i} \subseteq M|^{x_i, s_i} \subseteq M$, as desired. \square

Claim 21

Let $i \in \mathbb{N}$ and $x \in \mathcal{N}$. For each $(i, R(\bar{a})) \in \text{bf}_{i+1}(x)$, we have $R(x, \bar{a}) \in \text{async}_{\mathcal{P}}(D_i)$.

Proof

By definition of $\text{bf}_{i+1}(x)$, the pair $(i, R(\bar{a})) \in \text{bf}_{i+1}(x)$ implies that there are values $y \in \mathcal{N}$, $t \in \mathbb{N}$ and $u \in \mathbb{N}$ such that $\text{chosen}_R(y, t, x, u, \bar{a}) \in M$, $\text{glob}_M(y, t) = i$ and $\text{glob}_M(y, t) < i + 1 \leq \text{glob}_M(x, u)$. And $\text{glob}_M(y, t) = i$ gives us that $y = x_i$ and $t = s_i$. Thus $\text{chosen}_R(x_i, s_i, x, u, \bar{a}) \in M$.

All ground rules in G that can derive $\text{chosen}_R(x_i, s_i, x, u, \bar{a}) \in M$ are of the form (4), and hence $\text{cand}_R(x_i, s_i, x, u, \bar{a}) \in M$. Let $\psi \in G$ be an active ground rule

with head $\mathbf{cand}_R(x_i, s_i, x, u, \bar{a})$. Because $\psi \in G$, there is a rule $\varphi \in \text{pure}(\mathcal{P})$ and a valuation V so that ψ is obtained from φ by applying V and by subsequently removing all negative (ground) body literals, and so that $V(\text{neg}_\varphi) \cap M = \emptyset$. The rule φ is of the form (9), which implies that V must assign x_i and s_i respectively to the body location and timestamp variable that correspond to the sender, and that it must assign x and u respectively to the location and timestamp variable that correspond to the addressee. Let $\varphi' \in \mathcal{P}$ be the original asynchronous rule on which φ is based. Let φ'' be the corresponding rule in $\text{async}_{\mathcal{P}}$. From the construction of φ out of φ' and φ'' out of φ' , it follows that V can also be applied to φ'' . Note, $V(\text{head}_{\varphi''}) = R(x, \bar{a})$. We now show that V is satisfying for φ'' with respect to D_i , which causes $R(x, \bar{a}) \in \text{async}_{\mathcal{P}}(D_i)$, as desired. Specifically, we have to show $V(\text{pos}_{\varphi''}) \subseteq D_i$ and $V(\text{neg}_{\varphi''}) \cap D_i = \emptyset$.

First we show $V(\text{pos}_{\varphi''}) \subseteq D_i$. By construction of φ and φ'' , we have

$$\text{pos}_\psi|_{\text{sch}(\mathcal{P})^{\text{LT}}} = V(\text{pos}_\varphi)|_{\text{sch}(\mathcal{P})^{\text{LT}}} = V(\text{pos}_{\varphi''})^{\uparrow x_i, s_i}.$$

Since ψ is active, we have $\text{pos}_\psi|_{\text{sch}(\mathcal{P})^{\text{LT}}} \subseteq M$, and therefore $V(\text{pos}_{\varphi''})^{\uparrow x_i, s_i} \subseteq M$. Then, because the facts in $V(\text{pos}_{\varphi''})^{\uparrow x_i, s_i}$ are over $\text{sch}(\mathcal{P})^{\text{LT}}$ and have location specifier x_i and timestamp s_i , we can apply Claim 17 to know that $V(\text{pos}_{\varphi''}) \subseteq D_i$, as desired.

Now we show $V(\text{neg}_{\varphi''}) \cap D_i = \emptyset$. By construction of φ and φ'' , we have

$$V(\text{neg}_\varphi)|_{\text{sch}(\mathcal{P})^{\text{LT}}} = V(\text{neg}_{\varphi''})^{\uparrow x_i, s_i}.$$

By choice of φ and V , we have $V(\text{neg}_\varphi) \cap M = \emptyset$. Hence, $V(\text{neg}_{\varphi''})^{\uparrow x_i, s_i} \cap M = \emptyset$. Then, because the facts in $V(\text{neg}_{\varphi''})^{\uparrow x_i, s_i}$ are over $\text{sch}(\mathcal{P})^{\text{LT}}$ and have location specifier x_i and timestamp s_i , we can apply Claim 18 to know that $V(\text{neg}_{\varphi''}) \cap D_i = \emptyset$, as desired. \square

Claim 22

Let $i \in \mathbb{N}$. We have $M|_{x_i, s_i} = (D_i)^{\uparrow x_i, s_i}$. Intuitively, this means that the operational deductive fixpoint D_i during transition i , corresponding to step s_i of node x_i , is represented by M in an exact way.

Proof

Recall the notations from Section B.1. Let n denote the largest stratum number of the deductive rules of \mathcal{P} . We show by induction on $k = 0, 1, \dots, n$ that

$$(M^{\rightarrow k})|_{x_i, s_i} = (D_i^{\rightarrow k})^{\uparrow x_i, s_i}.$$

This will give us $(M^{\rightarrow n})|_{x_i, s_i} = (D_i^{\rightarrow n})^{\uparrow x_i, s_i} = (D_i)^{\uparrow x_i, s_i}$. Moreover, Claim 25 says that $(M^{\rightarrow n})|_{x_i, s_i} = M|_{x_i, s_i}$, and thus we obtain $M|_{x_i, s_i} = (D_i)^{\uparrow x_i, s_i}$, as desired.

Base case ($k = 0$) By definition,

$$M^{\rightarrow 0} = M^\blacktriangle \cup M^{\text{duc}, 0}.$$

But since there are no deductive ground rules in G with stratum 0, we have $M^{\text{duc},0} = \emptyset$. Hence,

$$\begin{aligned} (M^{\rightarrow 0})|_{x_i, s_i} &= (M^\blacktriangle)|_{x_i, s_i} \\ &= (M|_{\text{edb}(\mathcal{P})^{\text{LT}}})|_{x_i, s_i} \cup M^{\text{ind}}|_{x_i, s_i} \cup M^{\text{deliv}}|_{x_i, s_i}. \end{aligned} \quad (\text{B1})$$

Using Claim 23 and Claim 26, we can rewrite expression (B1) to the desired equality:

$$\begin{aligned} (M^{\rightarrow 0})|_{x_i, s_i} &= st_i(x_i)^{\uparrow x_i, s_i} \cup \text{untag}(m_i)^{\uparrow x_i, s_i} \\ &= (st_i(x_i) \cup \text{untag}(m_i))^{\uparrow x_i, s_i} \\ &= (D_i^{\rightarrow 0})^{\uparrow x_i, s_i}. \end{aligned}$$

Induction hypothesis For the induction hypothesis, we assume for a stratum number $k \geq 1$ that

$$(M^{\rightarrow k-1})|_{x_i, s_i} = (D_i^{\rightarrow k-1})^{\uparrow x_i, s_i}.$$

Inductive step We show that

$$(M^{\rightarrow k})|_{x_i, s_i} = (D_i^{\rightarrow k})^{\uparrow x_i, s_i}.$$

We show both inclusions separately, in Claims 27 and 28. \square

Claim 23

Let $i \in \mathbb{N}$. We have $st_i(x_i)^{\uparrow x_i, s_i} = (M|_{\text{edb}(\mathcal{P})^{\text{LT}}})|_{x_i, s_i} \cup M^{\text{ind}}|_{x_i, s_i}$.

Proof

By definition,

$$st_i(x_i) = ((M|_{\text{edb}(\mathcal{P})^{\text{LT}}})|_{x_i, s} \cup M^{\text{ind}}|_{x_i, s})^{\downarrow},$$

where $s = \text{loc}_M(i, x_i)$. Using Claim 15, we have $s = s_i$. Therefore,

$$st_i(x_i)^{\uparrow x_i, s_i} = (M|_{\text{edb}(\mathcal{P})^{\text{LT}}})|_{x_i, s_i} \cup M^{\text{ind}}|_{x_i, s_i}.$$

\square

Claim 24

For each fact $\text{cand}_R(x, s, y, u, \bar{a}) \in M$, there is a timestamp $u' \in \mathbb{N}$ such that $\text{chosen}_R(x, s, y, u', \bar{a}) \in M$, with possibly $u' = u$.

Proof

Towards a proof by contradiction, suppose there is no such timestamp u' . Now, because $\mathbf{cand}_R(x, s, y, u, \bar{a}) \in M$, the following ground rule, which is of the form (4), can not be in G , because otherwise $\mathbf{chosen}_R(x, s, y, u, \bar{a}) \in M$, which is assumed not to be possible:

$$\mathbf{chosen}_R(x, s, y, u, \bar{a}) \leftarrow \mathbf{cand}_R(x, s, y, u, \bar{a}).$$

Because rules of the form (4) contain a negative **other**...-atom in their body, the absence of the above ground rule from G implies $\mathbf{other}_R(x, s, y, u, \bar{a}) \in M$. This **other**_R-fact must be derived by a ground rule of the form (5):

$$\mathbf{other}_R(x, s, y, u, \bar{a}) \leftarrow \mathbf{cand}_R(x, s, y, u, \bar{a}), \mathbf{chosen}_R(x, s, y, u', \bar{a}), u \neq u'.$$

But this implies that $\mathbf{chosen}_R(x, s, y, u', \bar{a}) \in M$, which is a contradiction. \square

Claim 25

Let $i \in \mathbb{N}$. Let n denote the largest stratum number of the deductive rules of \mathcal{P} . We have $(M^{\rightarrow n})|_{x_i, s_i} = M|_{x_i, s_i}$.

Proof

First, since $M^{\rightarrow n} \subseteq M$, we immediately have $(M^{\rightarrow n})|_{x_i, s_i} \subseteq M|_{x_i, s_i}$.

Now, let $\mathbf{f} \in M|_{x_i, s_i}$. We show $\mathbf{f} \in (M^{\rightarrow n})|_{x_i, s_i}$. Since \mathbf{f} has location specifier x_i and timestamp s_i , we are left to show $\mathbf{f} \in M^{\rightarrow n}$. We have the following cases:

- Suppose $\mathbf{f} \in M|_{\text{edb}(\mathcal{P})^{\text{LT}}}$. Then $\mathbf{f} \in M^\bullet \subseteq M^{\rightarrow n}$.
- Suppose $\mathbf{f} \in M|_{\text{idb}(\mathcal{P})^{\text{LT}}}$. Then there is an active ground rule $\psi \in G$ with $\text{head}_\psi = \mathbf{f}$. As seen in Section B.1, rule ψ can be of three types: deductive, inductive and delivery. The last two cases would respectively imply $\mathbf{f} \in M^{\text{ind}}$ and $\mathbf{f} \in M^{\text{deliv}}$, giving $\mathbf{f} \in M^\bullet \subseteq M^{\rightarrow n}$. In the deductive case, rule ψ has a stratum number no larger than n , and hence $\mathbf{f} \in M^{\text{duc}, n} \subseteq M^{\rightarrow n}$.

\square

Claim 26

Let $i \in \mathbb{N}$. We have $M^{\text{deliv}}|_{x_i, s_i} = \text{untag}(m_i)^{\uparrow x_i, s_i}$.

Proof

Let $\mathbf{f} \in M^{\text{deliv}}|_{x_i, s_i}$. We show $\mathbf{f} \in \text{untag}(m_i)^{\uparrow x_i, s_i}$. Denote $\mathbf{f} = R(x_i, s_i, \bar{a})$. By definition of M^{deliv} , there is an active delivery rule $\psi \in G$ that derives \mathbf{f} :

$$R(x_i, s_i, \bar{a}) \leftarrow \mathbf{chosen}_R(y, t, x_i, s_i, \bar{a}).$$

Because this rule is active, we have $\mathbf{chosen}_R(y, t, x_i, s_i, \bar{a}) \in M$. Now, by definition of x_i and s_i , we have $\text{glob}_M(x_i, s_i) = i$. Hence, $(\text{glob}_M(y, t), R(\bar{a})) \in m_i$ and thus $R(\bar{a}) \in \text{untag}(m_i)$. Finally, we obtain $\mathbf{f} = R(x_i, s_i, \bar{a}) \in \text{untag}(m_i)^{\uparrow x_i, s_i}$, as desired.

Let $\mathbf{f} \in \text{untag}(m_i)^{\uparrow x_i, s_i}$. We show $\mathbf{f} \in M^{\text{deliv}}|_{x_i, s_i}$. Denote $\mathbf{f} = R(x_i, s_i, \bar{a})$. We

have $R(\bar{a}) \in \text{untag}(m_i)$. Thus, there is some tag $j \in \mathbb{N}$ such that $(j, R(\bar{a})) \in m_i$. By definition of m_i , there are values $y \in \mathcal{N}$, $t \in \mathbb{N}$, $z \in \mathcal{N}$ and $u \in \mathbb{N}$ such that

$$\text{chosen}_R(y, t, z, u, \bar{a}) \in M,$$

where $\text{glob}_M(y, t) = j$ and $\text{glob}_M(z, u) = i$. Here, $\text{glob}_M(z, u) = i$ implies $z = x_i$ and $u = s_i$. Hence, $\text{chosen}_R(y, t, x_i, s_i, \bar{a}) \in M$. Now, the following ground rule ψ is in G because (delivery) rules of the form (6) are always positive:

$$R(x_i, s_i, \bar{a}) \leftarrow \text{chosen}_R(y, t, x_i, s_i, \bar{a}).$$

This rule derives $\mathbf{f} = R(x_i, s_i, \bar{a}) \in M$ because its body-fact is in M . Hence, $\mathbf{f} \in M^{\text{deliv}}|_{x_i, s_i}$, as desired. \square

Claim 27

Let $i \in \mathbb{N}$. Let k be a stratum number (thus $k \geq 1$). Suppose that

$$(M^{\rightarrow k-1})|_{x_i, s_i} = (D_i^{\rightarrow k-1})^{\uparrow x_i, s_i}.$$

We have

$$(M^{\rightarrow k})|_{x_i, s_i} \subseteq (D_i^{\rightarrow k})^{\uparrow x_i, s_i}.$$

Proof

We consider the fixpoint computation of M , i.e., $M = \bigcup_{l \in \mathbb{N}} M_l$ with $M_0 = \text{decl}(H)$ and $M_l = T(M_{l-1})$ for each $l \geq 1$, where T is the immediate consequence operator of G . By the semantics of operator T , we have $M_{l-1} \subseteq M_l$.

We show by induction on $l = 0, 1, 2, \dots$, that

$$(M_l \cap M^{\rightarrow k})|_{x_i, s_i} \subseteq (D_i^{\rightarrow k})^{\uparrow x_i, s_i}.$$

This will imply that

$$\left(\left(\bigcup_{l \in \mathbb{N}} M_l \right) \cap M^{\rightarrow k} \right) |_{x_i, s_i} \subseteq (D_i^{\rightarrow k})^{\uparrow x_i, s_i}.$$

Hence, we obtain, as desired

$$(M \cap M^{\rightarrow k})|_{x_i, s_i} = (M^{\rightarrow k})|_{x_i, s_i} \subseteq (D_i^{\rightarrow k})^{\uparrow x_i, s_i}.$$

Before we start with the induction, recall from Section B.1 that

$$\begin{aligned} M^{\rightarrow k} &= M^\blacktriangle \cup M^{\text{duc}, k} \\ &= M|_{\text{edb}(\mathcal{P})}^{\text{LT}} \cup M^{\text{ind}} \cup M^{\text{deliv}} \cup M^{\text{duc}, k}. \end{aligned}$$

Base case ($l = 0$) We have $M_0 = \text{decl}(H)$. Thus M_0 contains no facts derived by deductive, inductive or delivery ground rules. Therefore,

$$M_0 \cap M^{\rightarrow k} = M|_{\text{edb}(\mathcal{P})^{\text{LT}}}.$$

Hence,

$$\begin{aligned} (M_0 \cap M^{\rightarrow k})|_{x_i, s_i} &\subseteq (M^\blacktriangle)|_{x_i, s_i} \\ &\subseteq (M^{\rightarrow k-1})|_{x_i, s_i}. \end{aligned}$$

And by using the given equality $(M^{\rightarrow k-1})|_{x_i, s_i} = (D_i^{\rightarrow k-1})^{\uparrow x_i, s_i}$, we obtain, as desired:

$$\begin{aligned} (M_0 \cap M^{\rightarrow k})|_{x_i, s_i} &\subseteq (D_i^{\rightarrow k-1})^{\uparrow x_i, s_i} \\ &\subseteq (D_i^{\rightarrow k})^{\uparrow x_i, s_i}. \end{aligned}$$

Induction hypothesis Let $l \geq 1$. We assume

$$(M_{l-1} \cap M^{\rightarrow k})|_{x_i, s_i} \subseteq (D_i^{\rightarrow k})^{\uparrow x_i, s_i}.$$

Inductive step We show

$$(M_l \cap M^{\rightarrow k})|_{x_i, s_i} \subseteq (D_i^{\rightarrow k})^{\uparrow x_i, s_i}.$$

Let $\mathbf{f} \in (M_l \cap M^{\rightarrow k})|_{x_i, s_i}$. If $\mathbf{f} \in M_{l-1}$ then $\mathbf{f} \in (M_{l-1} \cap M^{\rightarrow k})|_{x_i, s_i}$ and the induction hypothesis can be immediately applied. Now suppose that $\mathbf{f} \in M_l \setminus M_{l-1}$. Then there is a ground rule $\psi \in G$ with $\text{head}_\psi = \mathbf{f}$ that is active on M_{l-1} . We have $\text{pos}_\psi \subseteq M_{l-1}$. As we have seen in Section B.1, rule ψ can be of three types: deductive, inductive or a delivery. If ψ is an inductive rule or a delivery rule then

$$\begin{aligned} \mathbf{f} &\in M^{\text{ind}}|_{x_i, s_i} \cup M^{\text{deliv}}|_{x_i, s_i} \\ &\subseteq (M^\blacktriangle)|_{x_i, s_i} \subseteq (M^{\rightarrow k-1})|_{x_i, s_i} \\ &= (D_i^{\rightarrow k-1})^{\uparrow x_i, s_i} \subseteq (D_i^{\rightarrow k})^{\uparrow x_i, s_i}. \end{aligned}$$

Now suppose ψ is deductive. If ψ has stratum less than or equal to $k-1$, then $\mathbf{f} \in (M^{\rightarrow k-1})|_{x_i, s_i}$. In that case, the given equality $(M^{\rightarrow k-1})|_{x_i, s_i} = (D_i^{\rightarrow k-1})^{\uparrow x_i, s_i}$ gives $\mathbf{f} \in (D_i^{\rightarrow k-1})^{\uparrow x_i, s_i} \subseteq (D_i^{\rightarrow k})^{\uparrow x_i, s_i}$, as desired. Now suppose that ψ has stratum k . Because $\psi \in G$, there is a rule $\varphi \in \text{pure}(\mathcal{P})$ and valuation V so that ψ is obtained from φ by applying valuation V and subsequently removing the negative (ground) body literals, and so that $V(\text{neg}_\varphi) \cap M = \emptyset$. Let $\varphi' \in \mathcal{P}$ be the original deductive rule on which φ is based. Thus $\varphi' \in \text{deduc}_{\mathcal{P}}$ (see Section 5.1.2). By construction of φ out of φ' , valuation V can also be applied to rule φ' . We now show that V is satisfying for φ' during the computation of D_i , in stratum k . Since $V(\text{head}_\varphi) = \text{head}_\psi = \mathbf{f}$, this results in the derivation of $V(\text{head}_{\varphi'}) = \mathbf{f}^\downarrow \in D_i^{\rightarrow k}$ and thus $\mathbf{f} \in (D_i^{\rightarrow k})^{\uparrow x_i, s_i}$, as desired. It is sufficient to show $V(\text{pos}_{\varphi'}) \subseteq D_i^{\rightarrow k}$ and $V(\text{neg}_{\varphi'}) \cap D_i^{\rightarrow k-1} = \emptyset$ because by the syntactic stratification, if φ' uses relations positively then those relations are in stratum k or lower, and if φ' uses relations negatively then those relations are in a stratum strictly lower than k .

- We show $V(pos_{\varphi'}) \subseteq D_i^{\rightarrow k}$. First, by the relationship between φ and φ' , and because valuation V assigns x_i and s_i to respectively the body location variable and body timestamp variable of φ , we have $pos_{\psi} = V(pos_{\varphi}) = V(pos_{\varphi'})^{\uparrow x_i, s_i}$. By choice of ψ , we already know $pos_{\psi} \subseteq M_{l-1}$. If we could show $pos_{\psi} \subseteq M^{\rightarrow k}$ then $pos_{\psi} \subseteq (M_{l-1} \cap M^{\rightarrow k})|_{x_i, s_i}$, to which the induction hypothesis can be applied to obtain $pos_{\psi} = V(pos_{\varphi'})^{\uparrow x_i, s_i} \subseteq (D_i^{\rightarrow k})^{\uparrow x_i, s_i}$, resulting in $V(pos_{\varphi'}) \subseteq D_i^{\rightarrow k}$, as desired. Now we show $pos_{\psi} \subseteq M^{\rightarrow k}$. Let $\mathbf{g} \in pos_{\psi}$. If $\mathbf{g} \in M^{\blacktriangle}$ then we immediately have $\mathbf{g} \in M^{\rightarrow k}$. Now suppose that $\mathbf{g} \notin M^{\blacktriangle}$. Since $pos_{\psi} \subseteq M|_{x_i, s_i}$, we have $\mathbf{g} \in M|_{x_i, s_i} \setminus M^{\blacktriangle}$. Then Claim 25 implies there is an active deductive ground rule $\psi' \in G$ with $head_{\psi'} = \mathbf{g}$. But we are working with a syntactic stratification, and thus the stratum of ψ' can not be higher than the stratum of ψ , which is k . Hence $\mathbf{g} \in M^{\text{duc}, k} \subseteq M^{\rightarrow k}$.
- We show $V(neg_{\varphi'}) \cap D_i^{\rightarrow k-1} = \emptyset$. By choice of φ and V , we have $V(neg_{\varphi}) \cap M = \emptyset$. So,

$$V(neg_{\varphi}) \cap (M^{\rightarrow k-1})|_{x_i, s_i} = \emptyset.$$

By applying the given equality $(M^{\rightarrow k-1})|_{x_i, s_i} = (D_i^{\rightarrow k-1})^{\uparrow x_i, s_i}$, we then have $V(neg_{\varphi}) \cap (D_i^{\rightarrow k-1})^{\uparrow x_i, s_i} = \emptyset$. By the relationship between φ and φ' , we have $V(neg_{\varphi}) = V(neg_{\varphi'})^{\uparrow x_i, s_i}$. Thus $V(neg_{\varphi'}) \cap D_i^{\rightarrow k-1} = \emptyset$, as desired.

□

Claim 28

Let $i \in \mathbb{N}$. Let k be a stratum number (thus $k \geq 1$). Suppose that

$$(M^{\rightarrow k-1})|_{x_i, s_i} = (D_i^{\rightarrow k-1})^{\uparrow x_i, s_i}.$$

We have

$$(D_i^{\rightarrow k})^{\uparrow x_i, s_i} \subseteq (M^{\rightarrow k})|_{x_i, s_i}.$$

Proof

Recall that the semantics of stratum k in $deduc_{\mathcal{P}}$ is that of semi-positive Datalog⁻, with input $D_i^{\rightarrow k-1}$. So, we can consider $D_i^{\rightarrow k}$ to be a fixpoint, i.e., as the set $\bigcup_{l \in \mathbb{N}} A_l$ with $A_0 = D_i^{\rightarrow k-1}$ and $A_l = T(A_{l-1})$ for each $l \geq 1$, where T is the immediate consequence operator of stratum k in $deduc_{\mathcal{P}}$. We show by induction on $l = 0, 1, 2$, etc, that

$$(A_l)^{\uparrow x_i, s_i} \subseteq (M^{\rightarrow k})|_{x_i, s_i}.$$

This then gives us the desired result.

Base case ($l = 0$) We have $A_0 = D_i^{\rightarrow k-1}$. By applying the given equality, we obtain

$$(A_0)^{\uparrow x_i, s_i} = (D_i^{\rightarrow k-1})^{\uparrow x_i, s_i} = (M^{\rightarrow k-1})|_{x_i, s_i} \subseteq (M^{\rightarrow k})|_{x_i, s_i}.$$

Induction hypothesis Let $l \geq 1$. We assume

$$(A_{l-1})^{\uparrow x_i, s_i} \subseteq (M^{\rightarrow k})|_{x_i, s_i}.$$

Inductive step Let $\mathbf{f} \in A_l$. We show $\mathbf{f}^{\uparrow x_i, s_i} \in (M^{\rightarrow k})|_{x_i, s_i}$. If $\mathbf{f} \in A_{l-1}$ then the induction hypothesis can be applied to obtain the desired result. Now suppose $\mathbf{f} \in A_l \setminus A_{l-1}$. Let $\varphi \in \text{deduc}_{\mathcal{P}}$ and V be respectively a rule with stratum k and a valuation that together have derived $\mathbf{f} \in A_l$. Let $\varphi' \in \text{pure}(\mathcal{P})$ be the rule obtained from φ by applying transformation (1). Let V' be the extension of V to assign x_i and s_i respectively to the body location and timestamp variable of φ' , which are also both used in the head of φ' . Let ψ be the ground rule obtained from φ' by applying valuation V' and by subsequently removing all negative body literals. We show $\psi \in G$ and $\text{pos}_{\psi} \subseteq M$, which then implies

$$\text{head}_{\psi} = V'(\text{head}_{\varphi'}) = V(\text{head}_{\varphi})^{\uparrow x_i, s_i} = \mathbf{f}^{\uparrow x_i, s_i} \in M.$$

Moreover, because φ (and thus φ') has stratum k , rule ψ is an active deductive ground rule with stratum k , and thus $\mathbf{f}^{\uparrow x_i, s_i} \in (M^{\text{duc}, k})|_{x_i, s_i} \subseteq (M^{\rightarrow k})|_{x_i, s_i}$, as desired.

- To show $\psi \in G$, we require $V'(\text{neg}_{\varphi'}) \cap M = \emptyset$. Because V is satisfying for φ , and because negation is only applied to lower strata, we have

$$V(\text{neg}_{\varphi}) \cap D_i^{\rightarrow k-1} = \emptyset.$$

Thus

$$V(\text{neg}_{\varphi})^{\uparrow x_i, s_i} \cap (D_i^{\rightarrow k-1})^{\uparrow x_i, s_i} = \emptyset.$$

By the relationship between φ and φ' , we have $V(\text{neg}_{\varphi})^{\uparrow x_i, s_i} = V'(\text{neg}_{\varphi'})$, which gives us

$$V'(\text{neg}_{\varphi'}) \cap (D_i^{\rightarrow k-1})^{\uparrow x_i, s_i} = \emptyset.$$

And by using the given equality $(M^{\rightarrow k-1})|_{x_i, s_i} = (D_i^{\rightarrow k-1})^{\uparrow x_i, s_i}$, we have

$$V'(\text{neg}_{\varphi'}) \cap (M^{\rightarrow k-1})|_{x_i, s_i} = \emptyset.$$

Now, for the last step, we work towards a contradiction: suppose that there is a fact $\mathbf{g} \in V'(\text{neg}_{\varphi'}) \cap M$. From the construction of φ' , we know that \mathbf{g} is over $\text{sch}(\mathcal{P})^{\text{LT}}$ and has location specifier x_i and timestamp s_i .

- If \mathbf{g} is over $\text{edb}(\mathcal{P})^{\text{LT}}$ then $\mathbf{g} \in (M|_{\text{edb}(\mathcal{P})^{\text{LT}}})|_{x_i, s_i}$. Thus $\mathbf{g} \in (M^{\blacktriangle})|_{x_i, s_i} \subseteq (M^{\rightarrow k-1})|_{x_i, s_i}$, which is a contradiction.
- If \mathbf{g} is over $\text{idb}(\mathcal{P})^{\text{LT}}$ then there is an active ground rule $\psi' \in G$ with $\text{head}_{\psi'} = \mathbf{g}$. As seen in Section B.1, rule ψ' is either deductive, inductive or a delivery. The last two cases would imply that $\mathbf{g} \in (M^{\text{ind}} \cup M^{\text{deliv}})|_{x_i, s_i} \subseteq (M^{\blacktriangle})|_{x_i, s_i}$, which gives a contradiction like in the previous case. Now suppose that ψ' is deductive. Because the predicate of \mathbf{g} is used negatively in φ' and thus negatively in φ , the syntactic stratification assigns a smaller stratum number to ψ' than the stratum number of ψ , which is k . Hence, $\mathbf{g} \in (M^{\rightarrow k-1})|_{x_i, s_i}$, which is again a contradiction.

We conclude that $V'(\text{neg}_{\varphi'}) \cap M = \emptyset$.

- We show $pos_\psi \subseteq M$. Because V is satisfying for φ , we have

$$V(pos_\varphi) \subseteq A_{l-1}.$$

By the relationship between φ and φ' (and ψ), we have $V(pos_\varphi)^{\uparrow x_i, s_i} = V'(pos_{\varphi'}) = pos_\psi$. Thus

$$pos_\psi \subseteq (A_{l-1})^{\uparrow x_i, s_i}.$$

By now applying the induction hypothesis, we obtain, as desired:

$$pos_\psi \subseteq (M^{\rightarrow k})|^{x_i, s_i} \subseteq M.$$

□
